

cloudera[®]

Cloudera Release Guide

Important Notice

© 2010-2021 Cloudera, Inc. All rights reserved.

Cloudera, the Cloudera logo, and any other product or service names or slogans contained in this document are trademarks of Cloudera and its suppliers or licensors, and may not be copied, imitated or used, in whole or in part, without the prior written permission of Cloudera or the applicable trademark holder. If this documentation includes code, including but not limited to, code examples, Cloudera makes this available to you under the terms of the Apache License, Version 2.0, including any required notices. A copy of the Apache License Version 2.0, including any notices, is included herein. A copy of the Apache License Version 2.0 can also be found here: <https://opensource.org/licenses/Apache-2.0>

Hadoop and the Hadoop elephant logo are trademarks of the Apache Software Foundation. All other trademarks, registered trademarks, product names and company names or logos mentioned in this document are the property of their respective owners. Reference to any products, services, processes or other information, by trade name, trademark, manufacturer, supplier or otherwise does not constitute or imply endorsement, sponsorship or recommendation thereof by us.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Cloudera.

Cloudera may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Cloudera, the furnishing of this document does not give you any license to these patents, trademarks copyrights, or other intellectual property. For information about patents covering Cloudera products, see <http://tiny.cloudera.com/patents>.

The information in this document is subject to change without notice. Cloudera shall not be liable for any damages resulting from technical errors or omissions which may be present in this document, or from use of this document.

Cloudera, Inc.

**395 Page Mill Road
Palo Alto, CA 94306
info@cloudera.com
US: 1-888-789-1488
Intl: 1-650-362-0488
www.cloudera.com**

Release Information

Version: Cloudera Enterprise 5.x
Date: May 20, 2021

Table of Contents

Cloudera Enterprise 5.x Release Notes.....	19
---	-----------

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions.....	20
--	-----------

CDH Requirements for Cloudera Manager.....	20
CDH and Cloudera Manager Supported Operating Systems.....	20
<i>CDH and Cloudera Manager 5.16.x Supported Operating Systems.....</i>	<i>21</i>
<i>CDH and Cloudera Manager 5.15.x Supported Operating Systems.....</i>	<i>22</i>
<i>CDH and Cloudera Manager 5.14.x Supported Operating Systems.....</i>	<i>22</i>
<i>CDH and Cloudera Manager 5.13.x Supported Operating Systems.....</i>	<i>23</i>
<i>CDH and Cloudera Manager 5.12.x Supported Operating Systems.....</i>	<i>23</i>
<i>CDH and Cloudera Manager 5.11.x Supported Operating Systems.....</i>	<i>24</i>
<i>CDH and Cloudera Manager 5.10.x Supported Operating Systems.....</i>	<i>25</i>
<i>CDH and Cloudera Manager 5.9.x Supported Operating Systems.....</i>	<i>25</i>
<i>CDH and Cloudera Manager 5.8.x Supported Operating Systems.....</i>	<i>26</i>
<i>CDH and Cloudera Manager 5.7.x Supported Operating Systems.....</i>	<i>26</i>
<i>CDH and Cloudera Manager 5.6.x Supported Operating Systems.....</i>	<i>27</i>
<i>CDH and Cloudera Manager 5.5.x Supported Operating Systems.....</i>	<i>27</i>
<i>CDH and Cloudera Manager 5.4.x Supported Operating Systems.....</i>	<i>28</i>
<i>CDH and Cloudera Manager 5.3.x Supported Operating Systems.....</i>	<i>28</i>
<i>CDH and Cloudera Manager 5.2.x Supported Operating Systems.....</i>	<i>29</i>
<i>CDH and Cloudera Manager 5.1.x Supported Operating Systems.....</i>	<i>29</i>
<i>CDH and Cloudera Manager 5.0.x Supported Operating Systems.....</i>	<i>30</i>
Filesystem Requirements.....	30
Supported Filesystems.....	30
File Access Time.....	30
Filesystem Mount Options.....	31
CDH and Cloudera Manager Supported Databases.....	31
CDH and Cloudera Manager Supported JDK Versions.....	33
JDK 7.....	33
JDK 8.....	34
Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Requirement.....	35
Cloudera Manager Supported Browsers.....	35
Hue Supported Browsers.....	36
CDH and Cloudera Manager Supported Transport Layer Security Versions.....	36
Cloudera Manager Resource Requirements.....	38
CDH and Cloudera Manager Networking and Security Requirements.....	39
Product Compatibility Matrix for Apache Accumulo.....	42

- Product Compatibility Matrix for Apache Impala.....42
- Product Compatibility Matrix for CDK Powered By Apache Kafka.....44
- Kafka Client-Broker Compatibility across Kafka Versions.....45
- Kafka Client Versions Used by Apache Flume in CDH.....45
- Kafka Client Versions used by Apache Spark in CDH.....46
- Product Compatibility Matrix for Apache Kudu.....47
- Supported Operating Systems (Kudu 1.2.0 / CDH 5.10.x and higher).....47*
- Filesystem Requirements.....48*
- nscd Recommendation.....48*
- JDK Requirements.....48*
- Hardware Requirements.....48*
- Product Compatibility Matrix for Cloudera Navigator.....48
- Cloudera Navigator Supported Databases.....60*
- Cloudera Navigator Supported Browsers.....60*
- Cloudera Navigator Supported CDH and Managed Service Versions.....60*
- Data at Rest Encryption Requirements.....61
- Product Compatibility Matrix for Cloudera Navigator Encryption.....63
- Cloudera Navigator Key Trustee Server.....63*
- Key Trustee KMS.....65*
- Cloudera Navigator HSM KMS.....69*
- Cloudera Navigator Key HSM.....70*
- Cloudera Navigator Encrypt.....72*
- Product Compatibility Matrix for Apache Sentry.....74
- Product Compatibility Matrix for Apache Spark.....75
- Product Compatibility for Dell EMC Isilon.....76
- Product Compatibility Matrix for Backup and Disaster Recovery.....77
- Supported Replication Scenarios for Clusters using Isilon Storage.....80*
- Supported Configurations with Virtualization and Cloud Platforms.....80
- Amazon Web Services.....81*
- Google Cloud Platform.....81*
- Microsoft Azure.....81*
- VMware.....81*
- Red Hat OpenStack Platform (OSP) 11.....81*

Hardware Requirements Guide.....82

- Cloudera Manager.....82
- Cloudera Manager Server Storage Requirements.....82*
- Host Based Cloudera Manager Server Requirements.....82*
- Service Monitor Requirements.....82*
- Host Monitor.....83*
- Reports Manager.....84*
- Agent Hosts.....84*
- Event Server.....84*

<i>Alert Publisher</i>	84
Cloudera Navigator.....	84
Cloudera Data Science Workbench	86
CDH	87
<i>Accumulo</i>	87
<i>Flume</i>	88
<i>HDFS</i>	88
<i>HBase</i>	89
<i>Hive</i>	90
<i>HSM KMS</i>	91
<i>Hue</i>	92
<i>Impala</i>	92
<i>Kafka</i>	93
<i>Key Trustee Server</i>	93
<i>Key Trustee KMS</i>	94
<i>Kudu</i>	94
<i>Oozie</i>	94
<i>Search</i>	95
<i>Sentry</i>	95
<i>Spark</i>	96
<i>YARN</i>	96
<i>ZooKeeper</i>	98

Version and Download Information.....99

Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication.....	99
Cloudera Manager Version and Download Information.....	105
<i>Component Versions</i>	106
<i>Cloudera Manager 5.16.2.4505</i>	106
<i>Cloudera Manager 5.16.2</i>	108
<i>Previous Versions of Cloudera Manager 5</i>	110
CDH Version and Packaging Information.....	145
<i>CDH Download Information</i>	146
<i>CDH 5 Packaging and Tarball Information</i>	230
<i>Using the CDH 5 Maven Repository</i>	538
Cloudera Navigator Key Trustee Server Version and Download Information.....	1760
<i>Navigator Key Trustee Server 5.15</i>	1760
<i>Navigator Key Trustee Server 5.14</i>	1760
<i>Navigator Key Trustee Server 5.13</i>	1761
<i>Navigator Key Trustee Server 5.12</i>	1761
<i>Navigator Key Trustee Server 5.11</i>	1761
<i>Navigator Key Trustee Server 5.10</i>	1762
<i>Navigator Key Trustee Server 5.9</i>	1762
<i>Navigator Key Trustee Server 5.8</i>	1763

Navigator Key Trustee Server 5.7.....1763
Navigator Key Trustee Server 5.5.....1764
Navigator Key Trustee Server 5.4.....1764
Cloudera Navigator Key HSM Version and Download Information.....1766
Cloudera Navigator Key HSM 1.11.....1766
Cloudera Navigator Key HSM 1.10.....1766
Cloudera Navigator Key HSM 1.9.....1766
Cloudera Navigator Key HSM 1.8.....1767
Cloudera Navigator Key HSM 1.7.....1767
Cloudera Navigator Key HSM 1.6.....1767
Cloudera Navigator Key HSM 1.5.....1767
Cloudera Navigator Key HSM 1.4.....1768
Key Trustee KMS Version and Download Information.....1768
Key Trustee KMS 5.16.....1768
Key Trustee KMS 5.15.....1769
Key Trustee KMS 5.14.....1769
Key Trustee KMS 5.13.....1770
Key Trustee KMS 5.12.1.....1770
Key Trustee KMS 5.12.....1771
Key Trustee KMS 5.11.....1771
Key Trustee KMS 5.10.....1772
Key Trustee KMS 5.9.....1772
Key Trustee KMS 5.8.....1773
Key Trustee KMS 5.7.....1774
Key Trustee KMS 5.5.....1775
Key Trustee KMS 5.4.....1776
HSM KMS Version and Download Information.....1779
HSM KMS 5.15.....1779
HSM KMS 5.14.....1780
HSM KMS 5.13.....1780
HSM KMS 5.12.1.....1781
HSM KMS 5.12.....1781
Cloudera Navigator Encrypt Version and Download Information.....1782
Navigator Encrypt 3.16.....1782
Navigator Encrypt 3.15.....1783
Navigator Encrypt 3.14.....1784
Navigator Encrypt 3.13.....1785
Navigator Encrypt 3.12.....1785
Navigator Encrypt 3.11.....1786
Navigator Encrypt 3.10.....1787
Navigator Encrypt 3.9.....1788
Navigator Encrypt 3.8.....1789
Navigator Encrypt 3.7.....1790
Impala Version and Download Information.....1791

<i>Impala Download Information</i>	1792
--	------

What's New in Cloudera Documentation.....1800

What's New in Cloudera Documentation in May, 2019.....	1800
What's New in Cloudera Documentation in April, 2019.....	1800
What's New in Cloudera Documentation in March, 2019.....	1800
What's New in Cloudera Documentation in February, 2019.....	1801
What's New in Cloudera Documentation in January, 2019.....	1801
What's New in Cloudera Documentation in November, 2018.....	1801
What's New in Cloudera Documentation in October, 2018.....	1802
What's New in Cloudera Documentation in September, 2018.....	1803
What's New in Cloudera Documentation in August, 2018.....	1804
What's New in Cloudera Documentation in July, 2018.....	1804
What's New in Cloudera Documentation in June, 2018.....	1805
What's New in Cloudera Documentation in May, 2018.....	1806
What's New in Cloudera Documentation in April, 2018.....	1807
What's New in Cloudera Documentation in March, 2018.....	1807
What's New in Cloudera Documentation in February, 2018.....	1808
What's New in Cloudera Documentation in January, 2018.....	1808

CDH 5 Release Notes.....1812

New Features and Changes in CDH 5.....	1812
<i>What's New in CDH 5.16.x</i>	1812
<i>What's New in CDH 5.15.x</i>	1814
<i>What's New in CDH 5.14.x</i>	1816
<i>What's New in CDH 5.13.x</i>	1817
<i>What's New in CDH 5.12.x</i>	1819
<i>What's New in CDH 5.11.x</i>	1822
<i>What's New in CDH 5.10.x</i>	1824
<i>What's New In CDH 5.9.x</i>	1825
<i>What's New In CDH 5.8.x</i>	1827
<i>What's New In CDH 5.7.x</i>	1829
<i>What's New In CDH 5.6.x</i>	1834
<i>What's New In CDH 5.5.x</i>	1834
<i>What's New In CDH 5.4.x</i>	1838
<i>What's New In CDH 5.3.x</i>	1845
<i>What's New In CDH 5.2.x</i>	1848
<i>What's New in CDH 5.1.x</i>	1854
<i>What's New in CDH 5.0.x</i>	1858
<i>What's New In CDH 5 Beta Releases</i>	1861
<i>What's New in Apache Impala</i>	1869
<i>What's New in Apache Kudu</i>	1901

Incompatible Changes and Limitations.....	1916
<i>Apache Avro Incompatible Changes and Limitations.....</i>	<i>1916</i>
<i>Apache Crunch Incompatible Changes and Limitations.....</i>	<i>1916</i>
<i>Apache DataFu Incompatible Changes and Limitations.....</i>	<i>1916</i>
<i>Apache Flume Incompatible Changes and Limitations.....</i>	<i>1916</i>
<i>Apache Hadoop Incompatible Changes and Limitations.....</i>	<i>1916</i>
<i>Apache HBase Incompatible Changes and Limitations.....</i>	<i>1920</i>
<i>Apache Hive Incompatible Changes and Limitations.....</i>	<i>1924</i>
<i>Hue Incompatible Changes and Limitations.....</i>	<i>1925</i>
<i>Apache Impala Incompatible Changes and Limitations.....</i>	<i>1925</i>
<i>Cloudera Distribution of Apache Kafka Incompatible Changes and Limitations.....</i>	<i>1937</i>
<i>Kite Incompatible Changes and Limitations.....</i>	<i>1937</i>
<i>Apache Kudu Incompatible Changes and Limitations.....</i>	<i>1937</i>
<i>Llama Incompatible Changes and Limitations.....</i>	<i>1938</i>
<i>Apache Mahout Incompatible Changes and Limitations.....</i>	<i>1938</i>
<i>Apache Oozie Incompatible Changes and Limitations.....</i>	<i>1939</i>
<i>Apache Pig Incompatible Changes and Limitations.....</i>	<i>1939</i>
<i>Cloudera Search Incompatible Changes and Limitations.....</i>	<i>1939</i>
<i>Apache Sentry Incompatible Changes.....</i>	<i>1942</i>
<i>Apache Spark Incompatible Changes and Limitations.....</i>	<i>1942</i>
<i>Apache Sqoop Incompatible Changes and Limitations.....</i>	<i>1943</i>
<i>Apache Whirr Incompatible Changes and Limitations.....</i>	<i>1943</i>
<i>Apache ZooKeeper Incompatible Changes and Limitations.....</i>	<i>1943</i>
Known Issues in CDH 5.....	1944
<i>Operating System Known Issues.....</i>	<i>1944</i>
<i>Performance Known Issues.....</i>	<i>1949</i>
<i>Install and Upgrade Known Issues.....</i>	<i>1949</i>
<i>Apache Avro Known Issues.....</i>	<i>1951</i>
<i>Apache Flume Known Issues.....</i>	<i>1951</i>
<i>Apache Hadoop Known Issues.....</i>	<i>1952</i>
<i>Apache HBase Known Issues.....</i>	<i>1970</i>
<i>Apache Hive Known Issues.....</i>	<i>1979</i>
<i>Hue Known Issues.....</i>	<i>1991</i>
<i>Apache Impala Known Issues.....</i>	<i>1995</i>
<i>Cloudera Distribution of Apache Kafka Known Issues.....</i>	<i>2018</i>
<i>Apache Kudu Known Issues.....</i>	<i>2019</i>
<i>Apache Mahout Known Issues.....</i>	<i>2024</i>
<i>Apache Oozie Known Issues.....</i>	<i>2024</i>
<i>Apache Parquet Known Issues.....</i>	<i>2027</i>
<i>Apache Pig Known Issues.....</i>	<i>2027</i>
<i>Cloudera Search Known Issues.....</i>	<i>2028</i>
<i>Apache Sentry Known Issues.....</i>	<i>2035</i>
<i>Apache Spark Known Issues.....</i>	<i>2043</i>
<i>Apache Sqoop Known Issues.....</i>	<i>2049</i>
<i>Apache Phoenix Known Issues.....</i>	<i>2052</i>

<i>Apache ZooKeeper Known Issues</i>	2053
Issues Fixed in CDH 5.....	2054
<i>Issues Fixed in CDH 5.16.x</i>	2054
<i>Issues Fixed in CDH 5.15.x</i>	2063
<i>Issues Fixed in CDH 5.14.x</i>	2076
<i>Issues Fixed in CDH 5.13.x</i>	2106
<i>Issues Fixed in CDH 5.12.x</i>	2134
<i>Issues Fixed in CDH 5.11.x</i>	2151
<i>Issues Fixed in CDH 5.10.x</i>	2168
<i>Issues Fixed in CDH 5.9.x</i>	2194
<i>Issues Fixed in CDH 5.8.x</i>	2232
<i>Issues Fixed in CDH 5.7.x</i>	2264
<i>Issues Fixed in CDH 5.6.x</i>	2291
<i>Issues Fixed in CDH 5.5.x</i>	2297
<i>Issues Fixed in CDH 5.4.x</i>	2320
<i>Issues Fixed in CDH 5.3.x</i>	2346
<i>Issues Fixed in CDH 5.2.x</i>	2362
<i>Issues Fixed in CDH 5.1.x</i>	2370
<i>Issues Fixed in CDH 5.0.x</i>	2376
<i>Issues Fixed in CDH 5 Beta Releases</i>	2383
<i>Fixed Issues in Apache Impala</i>	2384
<i>Apache Kudu Fixed Issues</i>	2464

Cloudera Manager 5 Release Notes.....2472

New Features and Changes in Cloudera Manager 5.....	2472
<i>What's New in Cloudera Manager 5</i>	2472
<i>Incompatible Changes in Cloudera Manager 5</i>	2507
<i>Changed Features and Behaviors in Cloudera Manager 5</i>	2510
Known Issues and Workarounds in Cloudera Manager 5.....	2518
<i>Cloudera Manager 6.x issue with the service role Resume</i>	2518
<i>ZooKeeper JMX did not support TLS when managed by Cloudera Manager</i>	2519
<i>Yarn Application Masters fail when a concurrent Deploy Client Configuration job is executed</i>	2519
<i>Cloudera Manager Agent Failed to deactivate alternatives for parcels</i>	2520
<i>The client configuration deployment fails after upgrade</i>	2520
<i>Backup and Disaster Recovery replication with Isilon storage can fail</i>	2520
<i>Backup and Disaster Recovery replication can fail if one of the destination KMSs fail</i>	2520
<i>Cloudera Manager installation fails on MariaDB 10.2.8 and later</i>	2521
<i>Failures during Cloudera Manager installation or upgrade</i>	2521
<i>Upgrade to CDH 5.15.1 fails with CM 5.15.1 and OpenJDK</i>	2521
<i>After upgrade to Cloudera Manager 6, Kafka brokers and MirrorMaker marked as having stale configurations</i>	2521
<i>Parcel distribution fails and times out on Ubuntu 14.04 with Cloudera Express license</i>	2522
<i>Flume service with Kafka over TLS marked as stale after upgrading Cloudera Manager from 5.15.0</i>	2522
<i>Flume uses default values for Kafka TLS settings</i>	2522
<i>Upgrades from Cloudera Enterprise 5.15 to 6.0 Not Supported</i>	2523

<i>Cloudera Manager read-only user can access sensitive cluster information.....</i>	<i>2523</i>
<i>Accumulo health check fails on Cloudera Manager upgrade</i>	<i>2523</i>
<i>Cloudera Manager should not omit LDAP GROUP_MAPPING passwords from NodeManager.....</i>	<i>2523</i>
<i>Import step in Hive replication can fail in 5.15.0 due to a Hive concurrency issue.....</i>	<i>2524</i>
<i>Hard Restart of Cloudera Manager Agents May Cause Subsequent Service Errors.....</i>	<i>2524</i>
<i>Editing Cloudera Manager configurations in Google Chrome may lead to unexpected configuration loss.....</i>	<i>2524</i>
<i>Impala Dynamic Resource Pools wrongly gives everyone access to root pool (and all child pools).....</i>	<i>2524</i>
<i>Cross-site scripting vulnerability in Cloudera Manager</i>	<i>2524</i>
<i>Backup and Disaster Recovery (BDR) HDFS and Hive Replications may fail when replicating from source clusters which use HDFS HA and are managed by Cloudera Manager 5.13.3 or 5.14.2 to destination clusters managed by Cloudera 5.14.2.....</i>	<i>2525</i>
<i>Hive Replications Can Fail Intermittently.....</i>	<i>2526</i>
<i>HDFS DataNode stale configuration.....</i>	<i>2526</i>
<i>YARN MapReduce Job History stale configuration.....</i>	<i>2527</i>
<i>Logging issue slows down Backup and Disaster Recovery Hive and HDFS Replication jobs</i>	<i>2527</i>
<i>Cloudera Manager upgrade workflow incorrectly requires deploying some optional management roles</i>	<i>2527</i>
<i>ADLS Credentials in Log Files.....</i>	<i>2528</i>
<i>Installing Cloudera Manager on SLES Using the Installer Does Not Work.....</i>	<i>2528</i>
<i>Cloudera Navigator is not supported in multi-cluster deployments of Cloudera Manager versions 5.12.0, 5.12.1, and 5.13.0.....</i>	<i>2528</i>
<i>Backup and Disaster Recovery Hive replication fails intermittently.....</i>	<i>2528</i>
<i>Impala logs missing from diagnostic bundle.....</i>	<i>2529</i>
<i>Accumulo logs are not included in diagnostic bundles.....</i>	<i>2529</i>
<i>Backup and Disaster Recovery replication fails between two clusters when data is copied from one encryption zone to another when using CDH 5.12.0 or higher.....</i>	<i>2529</i>
<i>Backup and Disaster Recovery Hive replication fails during the Export Hive metastore step when using encrypted zones.....</i>	<i>2529</i>
<i>Backup and Disaster Recovery replication fails if the previous run is aborted.....</i>	<i>2529</i>
<i>Backup and Disaster Recovery (Hive) replication fails if the user directory is encrypted.....</i>	<i>2529</i>
<i>Upgrade to CDH 5.13 or higher Requires Pre-installation of Spark 2.1 or Spark 2.2.....</i>	<i>2530</i>
<i>Record Service no longer supported as of CDH 5.13; upgrade fails after from some lower versions.....</i>	<i>2530</i>
<i>Sentry may require increased Java heap settings before upgrading CDH to 5.13.....</i>	<i>2530</i>
<i>Renaming an empty directory on Amazon S3 when using DynamoDB with S3Guard may fail.....</i>	<i>2530</i>
<i>Change in Hue Load Balancer version causes Hue Load Balancer and Server to be marked with stale configuration.....</i>	<i>2531</i>
<i>Hue Load Balancer SSL Handshake error.....</i>	<i>2531</i>
<i>Apache MapReduce Jobs May Fail During Rolling Upgrade to CDH 5.11.0 or CDH 5.11.1.....</i>	<i>2531</i>
<i>Rolling Upgrade with GPL Extras Parcel Causes Oozie to Fail.....</i>	<i>2532</i>
<i>Maintenance State Minimal Block Replication staleness after upgrade.....</i>	<i>2532</i>
<i>YARN ACL configuration property staleness after upgrade.....</i>	<i>2532</i>
<i>Adding a service to an existing cluster.....</i>	<i>2532</i>
<i>Services die due to HDFS taking too long to start.....</i>	<i>2532</i>
<i>hostname parameter is not passed to Impala catalog role.....</i>	<i>2533</i>
<i>Zookeeper Package Installation fails with Debian 7</i>	<i>2533</i>
<i>Automated Cloudera Manager installer fails on Ubuntu 16.04.....</i>	<i>2534</i>

<i>Reboot of HDFS nodes running RHEL 7 or CentOS 7 shortly after enabling High Availability or migrating NameNodes can lead to data loss or other issues.....</i>	<i>2534</i>
<i>The Limit Nonsecure Container Executor Users property has no effect in Cloudera 5.11</i>	<i>2534</i>
<i>Graceful shutdown of Kafka brokers does not work as expected.....</i>	<i>2535</i>
<i>Spark Gateway roles should be added to every host.....</i>	<i>2535</i>
<i>HBase configuration is supplied for Hive when HBase service is not selected.....</i>	<i>2535</i>
<i>Hive Replication fails when Impala is SSL enabled but Hadoop services are not.....</i>	<i>2536</i>
<i>python-psycopg2 Dependency.....</i>	<i>2536</i>
<i>Pauses in Cloudera Manager after adding peer.....</i>	<i>2536</i>
<i>Cannot select time period in custom charts in C5.9.....</i>	<i>2537</i>
<i>Cloudera Manager 5.8.2 allows you to select nonexistent CDH 5.8.1 package installation.....</i>	<i>2537</i>
<i>Error when distributing parcels: No such torrent.....</i>	<i>2537</i>
<i>Hive Replication Metadata Transfer Step fails with Temporary AWS Credential Provider.....</i>	<i>2537</i>
<i>Hive table Views do not get restored from S3</i>	<i>2537</i>
<i>ACLs are not replicated when restoring Hive data from S3</i>	<i>2538</i>
<i>Snapshot diff is not working for Hive to S3 replication when data is deleted on source.....</i>	<i>2538</i>
<i>Block agents from heartbeating to a Cloudera Manager with different UUID until agent restart.....</i>	<i>2538</i>
<i>Cloudera Manager set catalogd default jvm memory to 4G can cause out of memory error on upgrade to Cloudera Manager 5.7 or higher.....</i>	<i>2538</i>
<i>Cloudera Manager 5.7.4 installer does not show Key Trustee KMS.....</i>	<i>2538</i>
<i>Class Not Found Error when upgrading to Cloudera Manager 5.7.2.....</i>	<i>2538</i>
<i>Kerberos setup fails on Debian 8.2.....</i>	<i>2538</i>
<i>Password in Cloudera Manager's <code>db.properties</code> file is not redacted.....</i>	<i>2539</i>
<i>Cluster provisioning fails.....</i>	<i>2539</i>
<i>Cloudera Manager can run out of memory if a remote repository URL is unreachable.....</i>	<i>2539</i>
<i>Clients can run Hive on Spark jobs even if Hive dependency on Spark is not configured.....</i>	<i>2539</i>
<i>The YARN NodeManager connectivity health test does not work for CDH 5.....</i>	<i>2540</i>
<i>HDFS HA clusters see NameNode failures when KDC connectivity is bad.....</i>	<i>2540</i>
<i>The HDFS File browser in Cloudera Manager fails when HDFS federation is enabled.....</i>	<i>2540</i>
<i>Hive Metastore canary fails to drop database.....</i>	<i>2540</i>
<i>Cloudera Manager upgrade fails due to incorrect Sqoop 2 path.....</i>	<i>2540</i>
<i>NameNode incorrectly reports missing blocks during rolling upgrade.....</i>	<i>2541</i>
<i>Using <code>ext3</code> for server dirs easily hit inode limit.....</i>	<i>2541</i>
<i>Backup and disaster recovery replication does not set MapReduce Java options</i>	<i>2541</i>
<i>Kafka 1.2 CSD conflicts with CSD included in Cloudera Manager.....</i>	<i>2541</i>
<i>Recommission host does not deploy client configurations.....</i>	<i>2541</i>
<i>Hive on Spark is not supported in Cloudera Manager and CDH 5.4 and CDH 5.5.....</i>	<i>2542</i>
<i>Upgrade wizard incorrectly upgrades the Sentry DB.....</i>	<i>2542</i>
<i>Cloudera Manager does not correctly generate client configurations for services deployed using CSDs.....</i>	<i>2542</i>
<i>Solr, Oozie and HttpFS fail when KMS and TLS/SSL are enabled using self-signed certificates.....</i>	<i>2542</i>
<i>Cloudera Manager 5.3.1 upgrade fails if Spark standalone and Kerberos are configured.....</i>	<i>2542</i>
<i>Adding Key Trustee KMS 5.4 to Cloudera Manager 5.5 displays warning.....</i>	<i>2543</i>
<i>KMS and Key Trustee ACLs do not work in Cloudera Manager.....</i>	<i>2543</i>
<i>Exporting and importing Hue database sometimes times out after 90 seconds.....</i>	<i>2543</i>
<i>Changing the Key Trustee Server hostname requires editing <code>keytrustee.conf</code>.....</i>	<i>2543</i>

<i>Hosts with Impala Llama roles must also have at least one YARN role.....</i>	<i>2543</i>
<i>The high availability wizard does not verify that there is a running ZooKeeper service.....</i>	<i>2544</i>
<i>Cloudera Manager Installation Path A fails on RHEL 5.7 due to PostgreSQL conflict.....</i>	<i>2544</i>
<i>Spurious warning on Accumulo 1.6 gateway hosts.....</i>	<i>2544</i>
<i>Accumulo 1.6 service log aggregation and search does not work.....</i>	<i>2544</i>
<i>Cloudera Manager incorrectly sizes Accumulo Tablet Server max heap size after 1.4.4-cdh4.5.0 to 1.6.0-cdh4.6.0 upgrade.....</i>	<i>2544</i>
<i>Accumulo installations using LZO do not indicate dependence on the GPL Extras parcel.....</i>	<i>2545</i>
<i>Created pools are not preserved when Dynamic Resource Pools page is used to configure YARN or Impala.....</i>	<i>2545</i>
<i>User should be prompted to add the AMON role when adding MapReduce to a CDH 5 cluster.....</i>	<i>2545</i>
<i>Enterprise license expiration alert not displayed until Cloudera Manager Server is restarted.....</i>	<i>2545</i>
<i>Configurations for decommissioned roles not migrated from MapReduce to YARN.....</i>	<i>2545</i>
<i>The HDFS command Roll Edits does not work in the UI when HDFS is federated.....</i>	<i>2545</i>
<i>Cloudera Manager reports a confusing version number if you have oozie-client, but not oozie installed on a CDH 4.4 node.....</i>	<i>2546</i>
<i>Cloudera Manager does not work with CDH 5.0.0 Beta 1.....</i>	<i>2546</i>
<i>On CDH 4.1 secure clusters managed by Cloudera Manager 4.8.1 and higher, the Impala Catalog server needs advanced configuration snippet update.....</i>	<i>2546</i>
<i>Rolling Upgrade to CDH 5 is not supported.....</i>	<i>2546</i>
<i>Error reading .zip file created with the Collect Diagnostic Data command.....</i>	<i>2546</i>
<i>After JobTracker failover, complete jobs from the previous active JobTracker are not visible.....</i>	<i>2546</i>
<i>After JobTracker failover, information about rerun jobs is not updated in Activity Monitor.....</i>	<i>2547</i>
<i>Installing on AWS, you must use private EC2 hostnames.....</i>	<i>2547</i>
<i>If HDFS uses Quorum-based Storage without HA enabled, the SecondaryNameNode cannot checkpoint.....</i>	<i>2547</i>
<i>Changing the rack configuration may temporarily cause mis-replicated blocks to be reported.....</i>	<i>2547</i>
<i>Cannot use '/' as a mount point with a Federated HDFS Nameservice.....</i>	<i>2547</i>
<i>Historical disk usage reports do not work with federated HDFS.....</i>	<i>2548</i>
<i>(CDH 4 only) Activity monitoring does not work on YARN activities.....</i>	<i>2548</i>
<i>HDFS monitoring configuration applies to all Nameservices.....</i>	<i>2548</i>
<i>Supported and Unsupported Replication Scenarios and Limitations.....</i>	<i>2548</i>
<i>Restoring snapshot of a file to an empty directory does not overwrite the directory.....</i>	<i>2548</i>
<i>HDFS Snapshot appears to fail if policy specifies duplicate directories.....</i>	<i>2548</i>
<i>Hive replication fails if "Force Overwrite" is not set.....</i>	<i>2548</i>
<i>Cloudera Manager set cataloged default JVM memory to 4G can cause an out of memory error during upgrade to Cloudera Manager 5.7 and higher.....</i>	<i>2549</i>
<i>Host Monitor Cannot Monitor NTP Status on SLES Hosts.....</i>	<i>2549</i>
<i>Issues Fixed in Cloudera Manager 5.....</i>	<i>2549</i>
<i>Issues Fixed in Cloudera Manager 5.16.2.....</i>	<i>2549</i>
<i>Issues Fixed in Cloudera Manager 5.16.1.....</i>	<i>2549</i>
<i>Issues Fixed in Cloudera Manager 5.15.2.....</i>	<i>2552</i>
<i>Issues Fixed in Cloudera Manager 5.15.1.....</i>	<i>2553</i>
<i>Issues Fixed in Cloudera Manager 5.15.0.....</i>	<i>2556</i>
<i>Issues Fixed in Cloudera Manager 5.14.4.....</i>	<i>2559</i>
<i>Issues Fixed in Cloudera Manager 5.14.3.....</i>	<i>2561</i>
<i>Issues Fixed in Cloudera Manager 5.14.2.....</i>	<i>2561</i>

Issues Fixed in Cloudera Manager 5.14.1.....2562
Issues Fixed in Cloudera Manager 5.14.0.....2563
Issues Fixed in Cloudera Manager 5.13.3.....2568
Issues Fixed in Cloudera Manager 5.13.2.....2568
Issues Fixed in Cloudera Manager 5.13.1.....2570
Issues Fixed in Cloudera Manager 5.13.0.....2573
Issues Fixed in Cloudera Manager 5.12.2.....2576
Issues Fixed in Cloudera Manager 5.12.1.....2578
Issues Fixed in Cloudera Manager 5.12.0.....2580
Issues Fixed in Cloudera Manager 5.11.2.....2585
Issues Fixed in Cloudera Manager 5.11.1.....2587
Issues Fixed in Cloudera Manager 5.11.0.....2589
Issues Fixed in Cloudera Manager 5.10.2.....2592
Issues Fixed in Cloudera Manager 5.10.1.....2593
Issues Fixed in Cloudera Manager 5.10.0.....2595
Issues Fixed in Cloudera Manager 5.9.3.....2597
Issues Fixed in Cloudera Manager 5.9.2.....2600
Issues Fixed in Cloudera Manager 5.9.1.....2601
Issues Fixed in Cloudera Manager 5.9.....2601
Issues Fixed in Cloudera Manager 5.8.5.....2606
Issues Fixed in Cloudera Manager 5.8.4.....2608
Issues Fixed in Cloudera Manager 5.8.3.....2609
Issues Fixed in Cloudera Manager 5.8.2.....2610
Issues Fixed in Cloudera Manager 5.8.1.....2612
Issues Fixed in Cloudera Manager 5.8.0.....2612
Issues Fixed in Cloudera Manager 5.7.6.....2621
Issues Fixed in Cloudera Manager 5.7.5.....2621
Issues Fixed in Cloudera Manager 5.7.4.....2622
Issues Fixed in Cloudera Manager 5.7.2.....2622
Issues Fixed in Cloudera Manager 5.7.1.....2624
Issues Fixed in Cloudera Manager 5.7.0.....2627
Issues Fixed in Cloudera Manager 5.6.2.....2630
Issues Fixed in Cloudera Manager 5.6.1.....2630
Issues Fixed in Cloudera Manager 5.6.0.....2631
Issues Fixed in Cloudera Manager 5.5.6.....2631
Issues Fixed in Cloudera Manager 5.5.4.....2632
Issues Fixed in Cloudera Manager 5.5.3.....2633
Issues Fixed in Cloudera Manager 5.5.2.....2634
Issues Fixed in Cloudera Manager 5.5.1.....2637
Issues Fixed in Cloudera Manager 5.5.0.....2638
Issues Fixed in Cloudera Manager 5.4.10.....2642
Issues Fixed in Cloudera Manager 5.4.9.....2643
Issues Fixed in Cloudera Manager 5.4.8.....2644
Issues Fixed in Cloudera Manager 5.4.7.....2644
Issues Fixed in Cloudera Manager 5.4.6.....2645

- Issues Fixed in Cloudera Manager 5.4.5*.....2646
- Issues Fixed in Cloudera Manager 5.4.3*.....2649
- Issues Fixed in Cloudera Manager 5.4.1*.....2655
- Issues Fixed in Cloudera Manager 5.4.0*.....2657
- Issues Fixed in Cloudera Manager 5.3.10*.....2658
- Issues Fixed in Cloudera Manager 5.3.9*.....2658
- Issues Fixed in Cloudera Manager 5.3.8*.....2659
- Issues Fixed in Cloudera Manager 5.3.7*.....2659
- Issues Fixed in Cloudera Manager 5.3.6*.....2660
- Issues Fixed in Cloudera Manager 5.3.4*.....2660
- Issues Fixed in Cloudera Manager 5.3.3*.....2661
- Issues Fixed in Cloudera Manager 5.3.2*.....2662
- Issues Fixed in Cloudera Manager 5.3.1*.....2664
- Issues Fixed in Cloudera Manager 5.3.0*.....2666
- Issues Fixed in Cloudera Manager 5.2.8*.....2667
- Issues Fixed in Cloudera Manager 5.2.7*.....2667
- Issues Fixed in Cloudera Manager 5.2.6*.....2668
- Issues Fixed in Cloudera Manager 5.2.5*.....2669
- Issues Fixed in Cloudera Manager 5.2.2*.....2669
- Issues Fixed in Cloudera Manager 5.2.1*.....2669
- Issues Fixed in Cloudera Manager 5.2.0*.....2670
- Issues Fixed in Cloudera Manager 5.1.7*.....2672
- Issues Fixed in Cloudera Manager 5.1.6*.....2672
- Fixed Issues in Cloudera Manager 5.1.5*.....2673
- Fixed Issues in Cloudera Manager 5.1.4*.....2673
- Issues Fixed in Cloudera Manager 5.1.3*.....2673
- Issues Fixed in Cloudera Manager 5.1.2*.....2674
- Issues Fixed in Cloudera Manager 5.1.1*.....2675
- Issues Fixed in Cloudera Manager 5.1.0*.....2675
- Fixed Issues in Cloudera Manager 5.0.8*.....2678
- Fixed Issues in Cloudera Manager 5.0.7*.....2678
- Fixed Issues in Cloudera Manager 5.0.6*.....2679
- Fixed Issues in Cloudera Manager 5.0.5*.....2679
- Issues Fixed in Cloudera Manager 5.0.2*.....2679
- Issues Fixed in Cloudera Manager 5.0.1*.....2679
- Issues Fixed in Cloudera Manager 5.0.0*.....2682
- Issues Fixed in Cloudera Manager 5.0.0 Beta 2*.....2685
- Issues Fixed in Cloudera Manager 5.0.0 Beta 1*.....2687

Cloudera Navigator Data Management Release Notes.....2688

- New Features and Changes in Cloudera Navigator Data Management*.....2688
- New Features in Cloudera Navigator*.....2688
- Changed Features in Cloudera Navigator*.....2709
- Known Issues and Workarounds in Cloudera Navigator Data Management*.....2714

<i>Access, Login, or Cloudera Navigator Console Issues</i>	2714
<i>Cloudera Manager Configuration Issues</i>	2716
<i>Multi-Cluster Environments</i>	2716
<i>AWS and Amazon S3</i>	2717
<i>Microsoft Azure</i>	2718
<i>Data Stewardship Dashboard</i>	2718
<i>Search and Managed Metadata</i>	2718
<i>Hive, Hue, Impala, Pig</i>	2719
<i>Navigator Audit Server</i>	2722
<i>Navigator Metadata Server</i>	2724
<i>Purge</i>	2729
<i>Spark</i>	2730
<i>Upgrade Issues and Limitations</i>	2731
<i>Issues Fixed in Cloudera Navigator Data Management</i>	2735
<i>Issues Fixed in Cloudera Navigator 2.15.2</i>	2735
<i>Issues Fixed in Cloudera Navigator 2.15.1</i>	2736
<i>Issues Fixed in Cloudera Navigator 2.14.2</i>	2740
<i>Issues Fixed in Cloudera Navigator 2.14.1</i>	2742
<i>Issues Fixed in Cloudera Navigator 2.14.0</i>	2744
<i>Issues Fixed in Cloudera Navigator 2.13.4</i>	2746
<i>Issues Fixed in Cloudera Navigator 2.13.3</i>	2748
<i>Issues Fixed in Cloudera Navigator 2.13.2</i>	2748
<i>Issues Fixed in Cloudera Navigator 2.13.1</i>	2749
<i>Issues Fixed in Cloudera Navigator 2.13.0</i>	2749
<i>Issues Fixed in Cloudera Navigator 2.12.3</i>	2751
<i>Issues Fixed in Cloudera Navigator 2.12.2</i>	2751
<i>Issues Fixed in Cloudera Navigator 2.12.1</i>	2751
<i>Issues Fixed in Cloudera Navigator 2.12.0</i>	2753
<i>Issues Fixed in Cloudera Navigator 2.11.2</i>	2754
<i>Issues Fixed in Cloudera Navigator 2.11.1</i>	2755
<i>Issues Fixed in Cloudera Navigator 2.11.0</i>	2755
<i>Issues Fixed in Cloudera Navigator 2.10.2</i>	2757
<i>Issues Fixed in Cloudera Navigator 2.10.1</i>	2758
<i>Issues Fixed in Cloudera Navigator 2.10.0</i>	2759
<i>Issues Fixed in Cloudera Navigator 2.9.2</i>	2760
<i>Issues Fixed in Cloudera Navigator 2.9.1</i>	2761
<i>Issues Fixed in Cloudera Navigator 2.9.0</i>	2762
<i>Issues Fixed in Cloudera Navigator 2.8.3</i>	2762
<i>Issues Fixed in Cloudera Navigator 2.8.1</i>	2763
<i>Issues Fixed in Cloudera Navigator 2.8.0</i>	2763
<i>Issues Fixed in Cloudera Navigator 2.7.3</i>	2764
<i>Issues Fixed in Cloudera Navigator 2.7.2</i>	2764
<i>Issues Fixed in Cloudera Navigator 2.7.1</i>	2765
<i>Issues Fixed in Cloudera Navigator 2.7.0</i>	2765
<i>Issues Fixed in Cloudera Navigator 2.6.5</i>	2768

<i>Issues Fixed in Cloudera Navigator 2.6.2</i>	2768
<i>Issues Fixed in Cloudera Navigator 2.6.1</i>	2769
<i>Issues Fixed in Cloudera Navigator 2.6.0</i>	2771
<i>Issues Fixed in Cloudera Navigator 2.5.0</i>	2772
<i>Issues Fixed in Cloudera Navigator 2.4.4</i>	2772
<i>Issues Fixed in Cloudera Navigator 2.4.2</i>	2772
<i>Issues Fixed in Cloudera Navigator 2.4.1</i>	2774
<i>Issues Fixed in Cloudera Navigator 2.4.0</i>	2774
<i>Issues Fixed in Cloudera Navigator 2.3.9</i>	2776
<i>Issues Fixed in Cloudera Navigator 2.3.8</i>	2777
<i>Issues Fixed in Cloudera Navigator 2.3.3</i>	2777
<i>Issues Fixed in Cloudera Navigator 2.3.1</i>	2778
<i>Issues Fixed in Cloudera Navigator 2.3.0</i>	2779
<i>Issues Fixed in Cloudera Navigator 2.2.9</i>	2779
<i>Issues Fixed in Cloudera Navigator 2.2.4</i>	2779
<i>Issues Fixed in Cloudera Navigator 2.2.3</i>	2779
<i>Issues Fixed in Cloudera Navigator 2.2.2</i>	2780
<i>Issues Fixed in Cloudera Navigator 2.2.1</i>	2780
<i>Issues Fixed in Cloudera Navigator 2.2.0</i>	2780
<i>Issues Fixed in Cloudera Navigator 2.1.6</i>	2780
<i>Issues Fixed in Cloudera Navigator 2.1.5</i>	2781
<i>Issues Fixed in Cloudera Navigator 2.1.4</i>	2781
<i>Issues Fixed in Cloudera Navigator 2.1.2</i>	2781
<i>Issues Fixed in Cloudera Navigator 2.1.1</i>	2782
<i>Issues Fixed in Cloudera Navigator 2.0.5</i>	2782
<i>Issues Fixed in Cloudera Navigator 2.0.3</i>	2782
<i>Issues Fixed in Cloudera Navigator 2.0.2</i>	2782
<i>Issues Fixed in Cloudera Navigator 2.0.1</i>	2782
<i>Issues Fixed in Cloudera Navigator 2.0.0</i>	2783
<i>Issues Fixed in Cloudera Navigator 1.2.0</i>	2784

Cloudera Navigator Key Trustee Server Release Notes.....2785

<i>New Features and Changes in Cloudera Navigator Key Trustee Server</i>	2785
<i>What's New in Cloudera Navigator Key Trustee Server</i>	2785
<i>Changed Features and Behaviors in Cloudera Navigator Key Trustee Server</i>	2786
<i>Known Issues and Workarounds in Cloudera Navigator Key Trustee Server</i>	2787
<i>Issues Fixed in Cloudera Navigator Key Trustee Server</i>	2790
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.15.0</i>	2790
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.14.0</i>	2791
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.13.0</i>	2791
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.12.0</i>	2791
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.11.1</i>	2791
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.11.0</i>	2791
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.10.1</i>	2791

<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.10.0</i>	2791
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.9.0</i>	2791
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.8.0</i>	2791
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.7.0</i>	2792
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.5.2</i>	2792
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.5.0</i>	2792
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.4.9</i>	2794
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.4.3</i>	2795
<i>Issues Fixed in Cloudera Navigator Key Trustee Server 5.4.0</i>	2796

Cloudera Navigator Key HSM Release Notes.....2797

New Features and Changes in Cloudera Navigator Key HSM.....	2797
<i>What's New in Cloudera Navigator Key HSM</i>	2797
Known Issues and Workarounds in Cloudera Navigator Key HSM.....	2798
Issues Fixed in Cloudera Navigator Key HSM.....	2799
<i>Issues Fixed in Cloudera Navigator Key HSM 1.11.0</i>	2799
<i>Issues Fixed in Cloudera Navigator Key HSM 1.10.0</i>	2799
<i>Issues Fixed in Cloudera Navigator Key HSM 1.9.0</i>	2799
<i>Issues Fixed in Cloudera Navigator Key HSM 1.8.0</i>	2799
<i>Issues Fixed in Cloudera Navigator Key HSM 1.7.0</i>	2799
<i>Issues Fixed in Cloudera Navigator Key HSM 1.6.0</i>	2799
<i>Issues Fixed in Cloudera Navigator Key HSM 1.5.1</i>	2799
<i>Issues Fixed in Cloudera Navigator Key HSM 1.5.0</i>	2799
<i>Issues Fixed in Cloudera Navigator Key HSM 1.4.0</i>	2800

Key Trustee KMS Release Notes.....2801

New Features in Key Trustee KMS.....	2801
<i>What's New in Key Trustee KMS</i>	2801
Known Issues and Workarounds in Key Trustee KMS.....	2803
Issues Fixed in Key Trustee KMS.....	2805
<i>Issues Fixed in Key Trustee KMS 5.16.0</i>	2805
<i>Issues Fixed in Key Trustee KMS 5.15.1</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.15.0</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.14.0</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.13.1</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.13.0</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.12.0</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.11.0</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.10.0</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.9.0</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.8.2</i>	2806
<i>Issues Fixed in Key Trustee KMS 5.8.0</i>	2807
<i>Issues Fixed in Key Trustee KMS 5.7.4</i>	2807

<i>Issues Fixed in Key Trustee KMS 5.7.1</i>	2807
<i>Issues Fixed in Key Trustee KMS 5.7.0</i>	2807
<i>Issues Fixed in Key Trustee KMS 5.5.4</i>	2807
<i>Issues Fixed in Key Trustee KMS 5.5.0</i>	2807
<i>Issues Fixed in Key Trustee KMS 5.4.3</i>	2808

Cloudera Navigator HSM KMS Release Notes.....2809

New Features and Changes in Cloudera Navigator HSM KMS.....	2809
<i>What's New in Cloudera Navigator HSM KMS</i>	2809
Known Issues and Workarounds in Cloudera Navigator HSM KMS.....	2809
Issues Fixed in Cloudera Navigator HSM KMS.....	2811
<i>Issues Fixed in Cloudera Navigator HSM KMS 5.15.0</i>	2811
<i>Issues Fixed in Cloudera Navigator HSM KMS 5.14.0</i>	2811
<i>Issues Fixed in Cloudera Navigator HSM KMS 5.13.0</i>	2812
<i>Issues Fixed in Cloudera HSM KMS 5.12.1</i>	2812
<i>Issues Fixed in Cloudera Navigator HSM KMS 5.12.0</i>	2812

Cloudera Navigator Encrypt Release Notes.....2813

New Features and Changes in Cloudera Navigator Encrypt.....	2813
<i>What's New in Cloudera Navigator Encrypt</i>	2813
<i>What's Changed in Cloudera Navigator Encrypt</i>	2814
Known Issues and Workarounds in Cloudera Navigator Encrypt.....	2814
Issues Fixed in Cloudera Navigator Encrypt.....	2816
<i>Issues Fixed in Cloudera Navigator Encrypt 3.16.0</i>	2816
<i>Issues Fixed in Cloudera Navigator Encrypt 3.15.0</i>	2816
<i>Issues Fixed in Cloudera Navigator Encrypt 3.14.0</i>	2817
<i>Issues Fixed in Cloudera Navigator Encrypt 3.13.0</i>	2817
<i>Issues Fixed in Cloudera Navigator Encrypt 3.12.2</i>	2817
<i>Issues Fixed in Cloudera Navigator Encrypt 3.12.0</i>	2817
<i>Issues Fixed in Cloudera Navigator Encrypt 3.11.0</i>	2817
<i>Issues Fixed in Cloudera Navigator Encrypt 3.10.0</i>	2817
<i>Issues Fixed in Cloudera Navigator Encrypt 3.9.0</i>	2817
<i>Issues Fixed in Cloudera Navigator Encrypt 3.8.0</i>	2818
<i>Issues Fixed in Cloudera Navigator Encrypt 3.7.1</i>	2818
<i>Issues Fixed in Cloudera Navigator Encrypt 3.7.0</i>	2819

Deprecated Items.....2820

Appendix: Apache License, Version 2.0.....2823

Cloudera Enterprise 5.x Release Notes

This guide contains release and download information for installers and administrators. It includes release notes as well as information about installation requirements, supported platforms, and version and download information. The guide also provides a release matrix that shows which major and minor release version of a product is supported with which release version of Cloudera Manager and CDH.



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

- [CDH 5 and Cloudera Manager 5 Requirements and Supported Versions](#) on page 20
- [Hardware Requirements Guide](#) on page 82
- [Version and Download Information](#) on page 99
- [What's New in Cloudera Documentation](#) on page 1800
- [CDH 5 Release Notes](#) on page 1812
- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Navigator Data Management Release Notes](#) on page 2688
- [Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785
- [Cloudera Navigator Key HSM Release Notes](#) on page 2797
- [Key Trustee KMS Release Notes](#) on page 2801
- [Cloudera Navigator Encrypt Release Notes](#) on page 2813
- [Deprecated Items](#) on page 2820

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

In an enterprise data hub, Cloudera Manager and CDH interact with several products such as Apache Accumulo, Apache Impala, Hue, Cloudera Search, and Cloudera Navigator. This guide provides information about which major and minor release version of a product is supported with which release version of CDH and Cloudera Manager.

Compatibility across different release versions of Cloudera Manager and CDH must be taken into account, especially when carrying out install/upgrade procedures.

JDK compatibility also varies across different Cloudera Manager and CDH versions. Certain versions of CDH 5 are compatible with both JDK 7 and JDK 8. In such cases, ensure all your services are deployed on the same major version. For example, you should not run Hadoop on JDK 7 while running Sqoop on JDK 8. Additionally, since Cloudera does not support mixed environments, all nodes in your cluster must be running the same major JDK version. Cloudera supports [Oracle](#) JDKs and, starting with Cloudera Enterprise 5.16, OpenJDK. For more information about the supported JDK versions, see [CDH and Cloudera Manager Supported JDK Versions](#) on page 33.

CDH Requirements for Cloudera Manager

Cloudera uses the following versioning convention: major.minor.maintenance. If a cluster runs Cloudera Manager 5.14.0, the major version is 5; the minor version is 14; and the maintenance version is 0.

The Cloudera Manager minor version must always be equal to or greater than the CDH minor version. Older versions of Cloudera Manager might not support features in newer versions of CDH.

For example, Cloudera Manager 5.12.0 can manage CDH 5.12.2 because the minor versions are equal. Cloudera Manager 5.12.0 cannot manage CDH 5.14.0 because the Cloudera Manager minor version, 12, is less than the CDH minor version, 14.



Important: Using Cloudera Manager 6.0.x to manage a CDH 5.15.x or CDH 5.16 cluster is not a supported configuration.

CDH and Cloudera Manager Supported Operating Systems



Note: Cloudera sometimes certifies OS releases between Cloudera Enterprise releases. When this occurs, the following tables are updated. This means that although an OS release might not be supported at the time of a Cloudera Enterprise release, it can be supported later.

CDH 5.x provides 64-bit packages for select versions of RHEL-compatible, SLES, Ubuntu, and Debian operating systems.



Important:

In order to be covered by Cloudera Support:

- All CDH hosts in a logical cluster must run on the same major OS release.
- Cloudera supports a *temporarily* mixed OS configuration during an OS upgrade project.
- Cloudera Manager must run on the same OS release as one of the CDH clusters it manages.

Cloudera recommends running the same minor release cross-cluster. However, the risk caused by running different minor OS releases is considered lower than the risk of running different major OS releases.

CDSW Gateway hosts may use RHEL/Centos 7.2, subject to some restrictions. See [Operating System Support for CDSW Gateway Hosts \(CDH 5.11 and higher only\)](#) on page 25

Points to note:

- Cloudera strongly discourages using RHEL 5 for new installations.
- RHEL / CentOS / OL 6.9 is supported in C5.7 and higher. In C5.7.x through 5.11.0, it has a [known cipher issue](#) (that is resolved in 5.11.1).
- Red Hat only supports specific upgrades from RHEL 6 to 7. Contact your OS vendor and review [What are the supported use cases for upgrading to RHEL 7?](#)
- SLES hosts running Cloudera Manager agents must use [SLES SDK 11 SP1](#).
- Cloudera does not support CDH cluster deployments in Docker containers.
- Cloudera is not responsible for developing or supporting SELinux policies. For more information, see [CDH and Cloudera Manager Networking and Security Requirements](#) on page 39.
- **Apache Kudu OS Requirements** - Kudu is supported only on a subset of the operating systems supported by CDH and Cloudera Manager. For details, refer [Supported Operating Systems \(Kudu 1.2.0 / CDH 5.10.x and higher\)](#) on page 47.
- **Cloudera Data Science Workbench OS Requirements** - Cloudera Data Science Workbench is supported only on a subset of the operating systems supported by CDH and Cloudera Manager. For details, refer the [CDSW documentation](#).

CDH and Cloudera Manager 5.16.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK	7.9 , 7.8, 7.7, 7.6, 7.5, 7.4, 7.3, 7.2, 7.1 6.10, 6.9, 6.8, 6.7, 6.6, 6.5, 6.4 5.11, 5.10, 5.7
Oracle Linux (OL)	7.5, 7.4, 7.3, 7.2, 7.1 (UEK default) 6.10, 6.9, 6.8 (UEK R2, R4) 6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11, 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	12 SP5 , 12 SP4*, 12 SP3, 12 SP2, 12 SP1 11 SP4, 11 SP3
Ubuntu/Debian	
Ubuntu	16.04 LTS (Xenial) 14.04 LTS (Trusty)
Debian	8.9, 8.4, 8.2 (Jessie) 7.8, 7.1, 7.0 (Wheezy)

* SLES 12 SP4 certification includes a known issue when MariaDB is being used as the Cloudera Manager server database: use MariaDB 10.0 instead of the default MariaDB 10.2 provided in the operating system release.

See [OPSAPS-52340](#).

CDH and Cloudera Manager 5.15.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK	7.9 , 7.8, 7.6, 7.5, 7.4, 7.3, 7.2, 7.1 6.10, 6.9, 6.8, 6.7, 6.6, 6.5, 6.4 5.11, 5.10, 5.7
Oracle Linux (OL)	7.5, 7.4, 7.3, 7.2, 7.1 (UEK default) 6.9, 6.8 (UEK R2, R4) 6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11, 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	12 SP3 , 12 SP2, 12 SP1 11 SP4, 11 SP3
Ubuntu/Debian	
Ubuntu	16.04 LTS (Xenial) 14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	8.9, 8.4, 8.2 (Jessie) 7.8, 7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.14.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK	7.6 , 7.5, 7.4, 7.3, 7.2, 7.1 6.10, 6.9, 6.8, 6.7, 6.6, 6.5, 6.4 5.11, 5.10, 5.7
Oracle Linux (OL)	7.5, 7.4, 7.3, 7.2, 7.1 (UEK default) 6.9, 6.8 (UEK R2, R4) 6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11, 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	12 SP2, 12 SP1 11 SP4, 11 SP3
Ubuntu/Debian	

Operating System	Version (bold=new)
Ubuntu	16.04 LTS (Xenial) 14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	8.9, 8.4, 8.2 (Jessie) 7.8, 7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.13.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 7.2</i>	7.6 , 7.5, 7.4, 7.3, 7.2, 7.1 6.10, 6.9, 6.8, 6.7, 6.6, 6.5, 6.4 5.11, 5.10, 5.7
Oracle Linux (OL)	7.5, 7.4, 7.3, 7.2, 7.1 (UEK default) 6.9, 6.8 (UEK R2, R4) 6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11, 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	12 SP2, 12 SP1 11 SP4, 11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	16.04 LTS (Xenial) 14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	8.9, 8.4, 8.2 (Jessie) 7.8, 7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.12.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 7.2</i>	7.5, 7.4, 7.3, 7.2, 7.1 6.10 , 6.9, 6.8, 6.7, 6.6, 6.5, 6.4 5.11, 5.10, 5.7
Oracle Linux (OL)	7.5, 7.4, 7.3, 7.2, 7.1 (UEK default) 6.9, 6.8 (UEK R2, R4)

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Operating System	Version (bold=new)
	6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11, 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	12 SP2, 12 SP1 11 SP4, 11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	16.04 LTS (Xenial) 14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	8.9 , 8.4, 8.2 (Jessie) 7.8, 7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.11.x Supported Operating Systems



Warning: RHEL/OL 6.x versions have a [known cipher issue](#) in 5.11.0 and resolved in 5.11.1.

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 7.2</i>	7.5, 7.4 , 7.3, 7.2, 7.1 6.9, 6.8, 6.7, 6.6, 6.5, 6.4 5.11, 5.10, 5.7
Oracle Linux (OL)	7.4 , 7.3, 7.2, 7.1 (UEK default) 6.9, 6.8 (UEK R2, R4) 6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11, 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	12 SP2 , 12 SP1 11 SP4, 11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	16.04 LTS (Xenial) 14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	8.4, 8.2 (Jessie)

Operating System	Version (bold=new)
	7.8, 7.1, 7.0 (Wheezy)

Operating System Support for CDSW Gateway Hosts (CDH 5.11 and higher only)

Cloudera Data Science Workbench is supported only on a subset of the operating systems supported by CDH and Cloudera Manager. For the complete list, see [CDSW - Supported Operating Systems](#).

A CDSW Gateway host that is dedicated to running Cloudera Data Science Workbench must use an operating system supported by CDSW even if the remaining hosts in your cluster are running any of the other [supported operating systems](#). All hosts must run the same version of a [supported JDK](#).

CDH and Cloudera Manager 5.10.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 6.7</i>	7.3 , 7.2, 7.1 6.9, 6.8, 6.7, 6.6, 6.5, 6.4 5.11, 5.10, 5.7
Oracle Linux (OL)	7.3 , 7.2, 7.1 (UEK default) 6.9, 6.8 (UEK R2, R4) 6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11, 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	12 SP1 11 SP4, 11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	8.4, 8.2 (Jessie) 7.8, 7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.9.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 6.7</i>	7.2, 7.1 6.9, 6.8 , 6.7, 6.6, 6.5, 6.4 5.11, 5.10, 5.7
Oracle Linux (OL)	7.2, 7.1 (UEK default) 6.9, 6.8 (UEK R2, R4)

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Operating System	Version (bold=new)
	6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11, 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	12 SP1 11 SP4, 11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	8.4 , 8.2 (Jessie) 7.8, 7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.8.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 6.6</i>	7.2, 7.1 6.9, 6.8, 6.7, 6.6, 6.5, 6.4 5.11, 5.10, 5.7
Oracle Linux (OL)	7.2, 7.1 (UEK default) 6.9 (UEK R2, R4) 6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11, 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	11 SP4, 11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	8.2 (Jessie) 7.8, 7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.7.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK	7.2 , 7.1

Operating System	Version (bold=new)
<i>Max SE Linux support: 6.6</i>	6.9, 6.8 , 6.7, 6.6, 6.5, 6.4 5.11 , 5.10, 5.7
Oracle Linux (OL)	7.2 , 7.1 (UEK default) 6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.11 , 5.10, 5.7 (UEK R2)
SUSE Linux Enterprise Server	
SLES	11 SP4 , 11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	7.8 , 7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.6.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 6.6</i>	7.1 6.7, 6.6, 6.5, 6.4 5.10, 5.7
Oracle Linux (OL)	7.1 (UEK default) 6.7, 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.10, 5.6 (UEK R2)
SUSE Linux Enterprise Server	
SLES	11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.5.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 6.6</i>	7.1 6.7 , 6.6, 6.5, 6.4

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Operating System	Version (bold=new)
	5.10, 5.7
Oracle Linux (OL)	7.1 (UEK default) 6.7 , 6.6, 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.10, 5.6 (UEK R2)
SUSE Linux Enterprise Server	
SLES	11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.4.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 6.5</i>	6.6 , 6.5, 6.4 5.10 , 5.7
Oracle Linux (OL)	6.6 , 6.5 (UEK R2, R3) 6.4 (UEK R2) 5.10 , 5.6 (UEK R2)
SUSE Linux Enterprise Server	
SLES	11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.3.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK <i>Max SE Linux support: 6.4</i>	6.5, 6.4 5.7
Oracle Linux (OL)	6.5 (UEK R2, R3) 6.4 (UEK R2) 5.6 (UEK R2)

Operating System	Version (bold=new)
SUSE Linux Enterprise Server	
SLES	11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.2.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK	6.5, 6.4
<i>Max SE Linux support: 6.4</i>	5.7
Oracle Linux (OL)	6.5 (UEK R2, R3) 6.4 (UEK R2) 5.6 (UEK R2)
SUSE Linux Enterprise Server	
SLES	11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	14.04 LTS (Trusty) 12.04 LTS (Precise)
Debian	7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.1.x Supported Operating Systems

Operating System	Version (bold=new)
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK	6.5, 6.4
<i>Max SE Linux support: 6.4</i>	5.7
Oracle Linux (OL)	6.5 (UEK R2, R3) 6.4 (UEK R2) 5.6 (UEK R2)
SUSE Linux Enterprise Server	
SLES	11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	12.04 LTS (Precise)

Operating System	Version (bold=new)
Debian	7.1, 7.0 (Wheezy)

CDH and Cloudera Manager 5.0.x Supported Operating Systems


Operating System	Version
Red Hat Enterprise Linux-compatible	
RHEL / CentOS / Oracle Linux RHCK	6.5, 6.4
<i>Max SE Linux support: 6.6</i>	5.7
Oracle Linux (OL)	6.5 (UEK R2, R3) 6.4 (UEK R2) 5.6 (UEK R2)
SUSE Linux Enterprise Server	
SLES	11 SP3, 11 SP2
Ubuntu/Debian	
Ubuntu	12.04 LTS (Precise)
Debian	7.1, 7.0 (Wheezy)

Filesystem Requirements

Supported Filesystems

The Hadoop Distributed File System (HDFS) is designed to run on top of an underlying filesystem in an operating system. Cloudera recommends that you use either of the following filesystems tested on the [supported operating systems](#):

- **ext3**: This is the most tested underlying filesystem for HDFS.
- **ext4**: This scalable extension of ext3 is supported in more recent Linux releases.



Important: Cloudera does not support in-place upgrades from ext3 to ext4. Cloudera recommends that you format disks as ext4 before using them as data directories.

- **XFS**: This is the default filesystem in RHEL 7.

Kudu Filesystem Requirements - Kudu is supported on ext4 and XFS. Kudu requires a kernel version and filesystem that supports hole punching. Hole punching is the use of the `fallocate(2)` system call with the `FALLOC_FL_PUNCH_HOLE` option set. For more details, see Kudu troubleshooting - [Error During Hole Punch Test](#).

File Access Time

Linux filesystems keep metadata that record when each file was accessed. This means that even reads result in a write to the disk. To speed up file reads, Cloudera recommends that you disable this option, called `atime`, using the mount option in `/etc/fstab`:

```
/dev/sdb1 /data1 ext4 defaults,noatime 0
```

Apply the change without rebooting:

```
mount -o remount /data1
```

Filesystem Mount Options

The filesystem `mount` options have a `sync` option that allows you to write synchronously.

However, using the `sync` option will lead to poor performance for services that write data to disks, such as HDFS, YARN, Kafka and Kudu. In CDH most writes are already replicated. Therefore, having synchronous writes to disks is unnecessary, expensive, and not worth the added safety it provides.

NFS and NAS options must not be used as DataNode Data Directory mounts, not even when using Hierarchical Storage features.

CDH and Cloudera Manager Supported Databases

Cloudera Manager and CDH come packaged with an embedded PostgreSQL database for use in a non-production environment. The embedded PostgreSQL database is not supported in production environments. For production environments, you must configure your cluster to use external databases.

In most cases (but not all), Cloudera supports versions of MariaDB, MySQL and PostgreSQL that are native to each supported Linux distribution.

After installing a database, upgrade to the latest patch and apply appropriate updates. Available updates may be specific to the operating system on which it is installed.

Notes:

- Use UTF8 encoding for all custom databases.
- For MySQL 5.6 and 5.7, you must install the *MySQL-shared-compat* or *MySQL-shared* package. This is required for the Cloudera Manager Agent package installation.
- MySQL GTID-based replication is not supported.
- Both the Community and Enterprise versions of MySQL are supported, as well as MySQL configured by the AWS RDS service.
- Cloudera supports Oracle Exadata and RAC instances when they serve as backend databases for CDH components without high availability. Cloudera software is architected with the assumption of a single database instance, and supports normal operations between Cloudera Enterprise and Oracle Exadata (or RAC) in such an environment. Cloudera cannot reasonably test and certify on each strategy for each RDBMS. Cloudera expects HA solutions for RDBMS to be transparent to Cloudera software, and therefore are not supported and debugged by Cloudera. It is the responsibility of the customer to provision, configure, and manage the RDBMS HA deployment, so that Cloudera software behaves exactly as it would when interfacing with a single, non-HA service. Cloudera will support and help customers troubleshoot issues when a cluster has HA enabled. While diagnosing database-related problems in Cloudera components, customers may be required to temporarily disable or bypass HA mechanisms for troubleshooting purposes. If an HA-related issue is found, it is the responsibility of the customer to engage with the database vendor so that a solution to that issue can be found.



Important: When you restart processes, the configuration for each of the services is redeployed using information saved in the Cloudera Manager database. If this information is not available, your cluster cannot start or function correctly. You must schedule and maintain regular backups of the Cloudera Manager database to recover the cluster in the event of the loss of this database. For more information, see [Backing Up Databases](#).

Table 1:  MySQL Support across CM/CDH 5 Releases

MySQL	Cloudera Manager/CDH 5.10 - 5.16	Cloudera Manager/CDH 5.1 - 5.9	Cloudera Manager/CDH 5.0
5.1	✓	✓	✓
5.5	✓	✓	✓
5.6	✓	✓	
5.7	✓		

Table 2:  MariaDB Support across CM/CDH 5 Releases

MariaDB	Cloudera Manager/CDH 5.9 - 5.16	Cloudera Manager/CDH 5.5 - 5.8
5.5	✓	✓
10.0	✓	
10.2	Recommend 10.0 because of the known issue OPSAPS-52340 .	

Table 3:  PostgreSQL Support across CM/CDH 5 Releases

PostgreSQL	Cloudera Manager/CDH 5.16	Cloudera Manager/CDH 5.15	Cloudera Manager/CDH 5.13 - 5.14	Cloudera Manager/CDH 5.11 - 5.12	Cloudera Manager/CDH 5.5 - 5.10	Cloudera Manager/CDH 5.2 - 5.4	Cloudera Manager/CDH 5.1	Cloudera Manager/CDH 5.0
8.1	✓	✓	✓	✓	✓	✓	✓	✓
8.3	✓	✓	✓	✓	✓	✓	✓	✓
8.4	✓	✓	✓	✓	✓	✓	✓	✓
9.1	✓	✓	✓	✓	✓	✓	✓	✓
9.2	✓	✓	✓	✓	✓	✓	✓	
9.3	✓	✓	✓	✓	✓	✓		
9.4	✓	✓	✓	✓	✓			
9.5	✓	✓	✓	✓				
9.6	✓	✓	✓					
10	✓	✓ See Warning below.						



Warning: Navigator Audit Server and Navigator Metadata Server released in Cloudera Manager 5.15.0 and 5.15.1 do not support PostgreSQL version 10. If you upgrade Cloudera Manager to use PostgreSQL version 10, make arrangements to provide an earlier version of PostgreSQL for Navigator Audit Server and Navigator Metadata Server.

Table 4:  Oracle Support across CM/CDH 5 Releases

Oracle	Cloudera Manager/CDH 5.13 - 5.16	Cloudera Manager/CDH 5.6 - 5.12	Cloudera Manager/CDH 5.0 - 5.5
11g R2	✓	✓	✓
12c R1	✓	✓	
12c R2	✓		
18	Cloudera Manager/CDH 5.16 only. Requires Java 8 JDBC driver		
19	Cloudera Manager/CDH 5.16 only. Requires Java 8 JDBC driver		

CDH and Cloudera Manager Supported JDK Versions



Note: A Java optimization called [compressed oops](#) (ordinary object pointers) enables a 64-bit JVM to address heap sizes up to about 32 GB using 4-byte pointers. For larger heap sizes, 8-byte pointers are required. This means that a heap size slightly less than 32 GB can hold more objects than a heap size slightly more than 32 GB.

If you do not need more than 32 GB heap, set your heap size to 31GB or less to avoid this issue. If you need 32 GB or more, set your heap size to 48 GB or higher to account for the larger pointers. In general, for heap sizes above 32 GB, multiply the amount of heap you need by 1.5.

Only 64 bit JDKs are supported. [Oracle JDK 7](#) is supported across all versions of Cloudera Manager 5 and CDH 5. [Oracle JDK 8](#) is supported in Cloudera Enterprise 5.3.x and higher.

OpenJDK 8 is supported in Cloudera Enterprise 5.16.1 and higher.

Oracle JDK 9 is not supported in any Cloudera Manager or CDH version.


Unless specifically excluded, Cloudera supports later updates to a major JDK release from the release that support was introduced. For example, Oracle JDK 8u102 was released prior to CDH 5.9, but is supported in CDH 5.3 and higher because support for Oracle JDK 8 was added in CDH 5.3. Cloudera excludes or removes support for select Java updates when security is jeopardized.

Running CDH nodes within the same cluster on different JDK releases is not supported. JDK release across a cluster needs to match the patch level.

JDK 7

All JDK 7 updates, from the minimum required version, are supported in Cloudera Manager/CDH 5.0 and higher unless specifically excluded. Updates above the minimum that are not listed are supported but not tested.

The Cloudera Manager repository contains **Oracle JDK 1.7.0_67** (for [example](#)) and can be automatically installed during a new installation or an upgrade.

 **Warning:**


- For CDH versions 5.13.0 and higher, Sentry requires JDK 7u76 or higher due to [JDK-8055949](#).
- CDK 3.0 and higher Powered By Apache Kafka require JDK 8, and do not support JDK 7.
- Versions 2.2 and higher of CDS Powered By Apache Spark, which can be installed on CDH 5, require JDK 8.
- OpenJDK 7 is not supported.

Table 5: Oracle JDK 7 versions that are tested and recommended


Oracle JDK Version	Notes
1.7u80	Recommended / Latest version tested
1.7u75	Recommended
1.7u67	Recommended
1.7u55	Minimum required

JDK 8

All Oracle JDK 8 updates, from the minimum required version, are supported in Cloudera Manager/CDH 5.3 and higher unless specifically excluded. Updates above the minimum that are not listed are supported but not tested.

 **Warning:**

- Oracle JDK 8u40, 8u45, and 8u60 are not supported due to JDK issues impacting CDH functionality:
 - Oracle JDK 8u40 and 8u45 are affected by [JDK-8077155](#), which affects HTTP authentication for certain web UIs.
 - Oracle JDK 8u60 is incompatible with the AWS SDK, and causes problem with DistCP. For more information, see the [KB article](#).
- OpenJDK 1.8u181 and higher are supported in Cloudera Enterprise 5.16.1 and higher.
- CDK 3.0 and higher Powered By Apache Kafka require JDK 8, and do not support JDK 7.
- JDK 8u75 is supported but has a Known Issue: [Oozie Web Console returns 500 error when Oozie server runs on JDK 8u75 or higher](#) on page 2026.
- Versions 2.2 and higher of CDS Powered By Apache Spark, which can be installed on CDH 5, require JDK 8.

 **Important:**

For JDK 8u241 and higher versions running on Kerberized clusters, you must disable referrals by setting `sun.security.krb5.disableReferrals=true`.

For example, with OpenJDK 1.8.0u242:

1. Open `/usr/lib/jvm/java-1.8.0-openjdk-1.8.0.242.b08-0.e17_7.x86_64/jre/lib/security/java.security` with a text editor.
2. Add `sun.security.krb5.disableReferrals=true` (it can be at the bottom of the file).
3. Add this property on each node that has the impacted JDK version.
4. Restart the applications using the JDK so the change takes effect.

For more information, see the [KB article](#).

Table 6: Oracle JDK 8 versions that are tested and recommended

Oracle JDK Version	Notes
1.8u181	Recommended / Latest version tested
1.8u162	Recommended / Latest version tested
1.8u144	Recommended
1.8u131	Recommended
1.8u121	Recommended
1.8u111	Recommended
1.8u102	Recommended
1.8u91	Recommended
1.8u74	Recommended
1.8u31	Minimum required

Table 7: OpenJDK 1.8 versions that are tested and recommended

OpenJDK Version	Notes
1.8u181	Minimum required / Latest version tested

Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Requirement

If you are using CentOS/Red Hat Enterprise Linux 5.6 or higher, or Ubuntu, you must install the [Java Cryptography Extension \(JCE\) Unlimited Strength Jurisdiction Policy File](#) on all cluster hosts (including gateway hosts). This is to ensure that JDK uses the same default encryption type (`aes256-cts`) as that used by Red Hat/CentOS operating systems, Kerberos, and the rest of the CDH cluster.

On a managed deployment, Cloudera Manager can install the JCE on hosts. For unmanaged deployments, use the installation instructions in the `README.txt` file included in `jce_policy-x.zip`.

- [JCE for Java SE 7](#)
- For JDK 1.8.0_150 (and lower) - [JCE for Java SE 8](#)

Enabling Unlimited Strength Encryption for JDK 1.8.0_151 (and later)

- As of **JDK 1.8.0_151**, unlimited strength encryption can be enabled using the `java.security` file as documented in the [JDK 1.8.0_151 release notes](#). You do not need to install the JCE policy files.
- As of **JDK 1.8.0_161**, unlimited strength encryption has been [enabled by default](#). No further action is required.

Cloudera Manager Supported Browsers

The Cloudera Manager Admin Console, which you use to install, configure, manage, and monitor services, supports the latest version of the following browsers:


- Mozilla Firefox
- Google Chrome
- Internet Explorer
- Safari.

Hue Supported Browsers

Hue works with the two most recent [LTS](#) (long term support) or [ESR](#) (extended support release) browsers. Cookies and JavaScript must be on.

- **Chrome:** [Version history](#)
- **Firefox:** [Version history](#)
- **Internet Explorer:** [Version history](#)
- **Safari (Mac only):** [Version history](#)
- **Windows Edge**[Version history](#)

Hue can display in older, and other, browsers, but you might not have access to all of its features.

 **Important:** To see all icons in the Hue Web UI, users with IE and HTTPS must add a Load Balancer.

CDH and Cloudera Manager Supported Transport Layer Security Versions

The following components are supported by the indicated versions of Transport Layer Security (TLS):

Table 8: Components Supported by TLS

Component	Role	Name	Port	Version
Cloudera Manager	Cloudera Manager Server		7182	TLS 1.2
Cloudera Manager	Cloudera Manager Server		7183	TLS 1.2
			9099	TLS 1.2
Flume		Avro Source/Sink		TLS 1.2
Flume		Flume HTTP Source/Sink		TLS 1.2
HBase	Master	HBase Master Web UI Port	60010	TLS 1.2
HDFS	NameNode	Secure NameNode Web UI Port	50470	TLS 1.2
HDFS	Secondary NameNode	Secure Secondary NameNode Web UI Port	50495	TLS 1.2
HDFS	HttpFS	REST Port	14000	TLS 1.1, TLS 1.2
Hive	HiveServer2	HiveServer2 Port	10000	TLS 1.2
Hue	Hue Server	Hue HTTP Port	8888	TLS 1.2
Impala	Impala Daemon	Impala Daemon Beeswax Port	21000	TLS 1.0, TLS 1.1, TLS 1.2 We recommend that clients use the highest supported version, TLS 1.2.
Impala	Impala Daemon	Impala Daemon HiveServer2 Port	21050	TLS 1.0, TLS 1.1, TLS 1.2 We recommend that clients use the highest

Component	Role	Name	Port	Version
				supported version, TLS 1.2.
Impala	Impala Daemon	Impala Daemon Backend Port	22000	TLS 1.0, TLS 1.1, TLS 1.2 We recommend that clients use the highest supported version, TLS 1.2.
Impala	Impala StateStore	StateStore Service Port	24000	TLS 1.0, TLS 1.1, TLS 1.2 We recommend that clients use the highest supported version, TLS 1.2.
Impala	Impala Daemon	Impala Daemon HTTP Server Port	25000	TLS 1.0, TLS 1.1, TLS 1.2 We recommend that clients use the highest supported version, TLS 1.2.
Impala	Impala StateStore	StateStore HTTP Server Port	25010	TLS 1.0, TLS 1.1, TLS 1.2 We recommend that clients use the highest supported version, TLS 1.2.
Impala	Impala Catalog Server	Catalog Server HTTP Server Port	25020	TLS 1.0, TLS 1.1, TLS 1.2 We recommend that clients use the highest supported version, TLS 1.2.
Impala	Impala Catalog Server	Catalog Server Service Port	26000	TLS 1.0, TLS 1.1, TLS 1.2 We recommend that clients use the highest supported version, TLS 1.2.
Oozie	Oozie Server	Oozie HTTPS Port	11443	TLS 1.1, TLS 1.2
Solr	Solr Server	Solr HTTP Port	8983	TLS 1.1, TLS 1.2
Solr	Solr Server	Solr HTTPS Port	8985	TLS 1.1, TLS 1.2
Spark	History Server		18080	TLS 1.2
YARN	ResourceManager	ResourceManager Web Application HTTP Port	8090	TLS 1.2
YARN	JobHistory Server	MRv1 JobHistory Web Application HTTP Port	19890	TLS 1.2

Cloudera Manager Resource Requirements

Cloudera Manager requires the following resources:

- **Disk Space**
 - **Cloudera Manager Server**
 - 5 GB on the partition hosting `/var`.
 - 500 MB on the partition hosting `/usr`.
 - For parcels, the space required depends on the number of parcels you download to the Cloudera Manager Server and distribute to Agent hosts. You can download multiple parcels of the same product, of different versions and different builds. If you are managing multiple clusters, only one parcel of a product/version/build/distribution is downloaded on the Cloudera Manager Server—not one per cluster. In the local parcel repository on the Cloudera Manager Server, the approximate sizes of the various parcels are as follows:
 - CDH 5 (which includes Impala and Search) - 1.5 GB per parcel (packed), 2 GB per parcel (unpacked)
 - Impala - 200 MB per parcel
 - Cloudera Search - 400 MB per parcel
 - **Cloudera Management Service** -The Host Monitor and Service Monitor databases are stored on the partition hosting `/var`. Ensure that you have at least 20 GB available on this partition.
 - **Agents** - On Agent hosts, each unpacked parcel requires about three times the space of the downloaded parcel on the Cloudera Manager Server. By default, unpacked parcels are located in `/opt/cloudera/parcels`.
- **RAM** - 4 GB is recommended for most cases and is required when using Oracle databases. 2 GB might be sufficient for non-Oracle deployments with fewer than 100 hosts. However, to run the Cloudera Manager Server on a machine with 2 GB of RAM, you must tune down its maximum heap size (by modifying `-Xmx` in `/etc/default/cloudera-scm-server`). Otherwise the kernel might kill the Server for consuming too much RAM.
- **Python** - Cloudera Manager requires Python 2.4 or higher (but is not compatible with Python 3.0 or higher). Hue in CDH 5 and package installs of CDH 5 require Python 2.6 or 2.7. All supported operating systems include Python version 2.4 or higher. Cloudera Manager is compatible with Python 2.4 through the latest version of Python 2.x. Cloudera Manager does not support Python 3.0 and higher.
- **Perl** - Cloudera Manager requires [perl](#).
- **python-psycopg2 package** - Cloudera Manager 5.8 and higher has a new dependency on the package `python-psycopg2`. This package is not available in standard SLES 11 and SLES 12 repositories. You need to add this repository or install it manually to any machine that runs the Cloudera Manager Agent before you install or upgrade Cloudera Manager.

If the Cloudera Manager Server and Agent run on the same host, install the Cloudera Manager Server first and then add the `python-psycopg2` repository or package. After adding the repository or package, install the Cloudera Manager Agent.

Download the `python-psycopg2` repository or package from the following URL by selecting the correct SLES version: <https://software.opensuse.org/package/python-psycopg2>.

- **iproute package** - Cloudera Manager 5.12 and higher has a new dependency on the `iproute` package. Any machine that runs the Cloudera Manager Agent requires the package. Which version you need depends on the operating system:

Table 9: iproute package

Operating System	iproute version
CentOS 6.x	iproute-2.6
CentOS 7.x	iproute-3.10

Operating System	iproute version
RHEL 6.x	iproute-2.6
RHEL 7.x	iproute-3.10
SLES 11 SP3	iproute2-2.6
SLES 11 SP4	iproute2-3.0
SLES 12 and 12 SP1	iproute2-3.12
SLES 12 SP4	iproute2-4.4

CDH and Cloudera Manager Networking and Security Requirements

The hosts in a Cloudera Manager deployment must satisfy the following networking and security requirements:

- **Networking Protocols Support**

CDH requires IPv4. IPv6 is not supported and must be disabled.



Note: Contact your OS vendor for help disabling IPv6.

See also [Configuring Network Names](#).

- **Multihoming Support**

– Multihoming CDH or Cloudera Manager is not supported outside specifically certified Cloudera partner appliances. Cloudera finds that current Hadoop architectures combined with modern network infrastructures and security practices remove the need for multihoming. Multihoming, however, is beneficial internally in appliance form factors to take advantage of high-bandwidth InfiniBand interconnects.

Although some subareas of the product may work with unsupported custom multihoming configurations, there are known issues with multihoming. In addition, unknown issues may arise because multihoming is not covered by our test matrix outside the Cloudera-certified partner appliances.

- **Entropy**

Data at rest encryption requires sufficient [entropy](#) to ensure randomness.

See entropy requirements in [Data at Rest Encryption Requirements](#) on page 61.

- Cluster hosts must have a working network name resolution system and correctly formatted `/etc/hosts` file. All cluster hosts must have properly configured forward and reverse host resolution through DNS. The `/etc/hosts` files must:

- Contain consistent information about hostnames and IP addresses across all hosts
- Not contain uppercase hostnames
- Not contain duplicate IP addresses

Cluster hosts must not use aliases, either in `/etc/hosts` or in configuring DNS. A properly formatted `/etc/hosts` file should be similar to the following example:

```
127.0.0.1 localhost.localdomain localhost
192.168.1.1 cluster-01.example.com cluster-01
192.168.1.2 cluster-02.example.com cluster-02
192.168.1.3 cluster-03.example.com cluster-03
```

- In most cases, the Cloudera Manager Server must have SSH access to the cluster hosts when you run the installation or upgrade wizard. You must log in using a root account or an account that has password-less sudo permission.

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

For authentication during the installation and upgrade procedures, you must either enter the password or upload a public and private key pair for the root or sudo user account. If you want to use a public and private key pair, the public key must be installed on the cluster hosts before you use Cloudera Manager.

Cloudera Manager uses SSH only during the initial install or upgrade. Once the cluster is set up, you can disable root SSH access or change the root password. Cloudera Manager does not save SSH credentials, and all credential information is discarded when the installation is complete.

- If [single user mode](#) is not enabled, the Cloudera Manager Agent runs as root so that it can make sure the required directories are created and that processes and files are owned by the appropriate user (for example, the `hdfs` and `mapred` users).
- No blocking is done by Security-Enhanced Linux (SELinux).



Note: Cloudera Enterprise, with the exception of Cloudera Navigator Encrypt, is supported on platforms with Security-Enhanced Linux (SELinux) enabled and in `enforcing` mode. Cloudera is not responsible for SELinux policy development, support, or enforcement. If you experience issues running Cloudera software with SELinux enabled, contact your OS provider for assistance.

If you are using SELinux in `enforcing` mode, Cloudera Support can request that you disable SELinux or change the mode to `permissive` to rule out SELinux as a factor when investigating reported issues.

- No blocking by iptables or firewalls; port 7180 must be open because it is used to access Cloudera Manager after installation. Cloudera Manager communicates using specific ports, which must be open.
- For RHEL and CentOS, the `/etc/sysconfig/network` file on each host must contain the hostname you have just set (or verified) for that host.
- Cloudera Manager and CDH use several user accounts and groups to complete their tasks. The set of user accounts and groups varies according to the components you choose to install. Do not delete these accounts or groups and do not modify their permissions and rights. Ensure that no existing systems prevent these accounts and groups from functioning. For example, if you have scripts that delete user accounts not in a whitelist, add these accounts to the list of permitted accounts. Cloudera Manager, CDH, and managed services create and use the following accounts and groups:

Table 10: Users and Groups

Component (Version)	Unix User ID	Groups	Functionality
Cloudera Manager (all versions)	<code>cloudera-scm</code>	<code>cloudera-scm</code>	Clusters managed by Cloudera Manager run Cloudera Manager Server, monitoring roles, and other Cloudera Server processes as <code>cloudera-scm</code> . Requires keytab file named <code>cmf.keytab</code> because name is hard-coded in Cloudera Manager.
Apache Accumulo (Accumulo 1.4.3 and higher)	<code>accumulo</code>	<code>accumulo</code>	Accumulo processes run as this user.
Apache Avro	<code>~</code>	<code>~</code>	No special user:group.
Apache Flume	<code>flume</code>	<code>flume</code>	The sink that writes to HDFS as user must have write privileges.
Apache HBase	<code>hbase</code>	<code>hbase</code>	The Master and the RegionServer processes run as this user.

Component (Version)	Unix User ID	Groups	Functionality
HDFS	hdfs	hdfs, hadoop	The NameNode and DataNodes run as this user, and the HDFS root directory as well as the directories used for edit logs should be owned by it.
Apache Hive	hive	hive	The HiveServer2 process and the Hive Metastore processes run as this user. A user must be defined for Hive access to its Metastore DB (for example, MySQL or Postgres) but it can be any identifier and does not correspond to a Unix uid. This is <code>javax.jdo.option.ConnectionUserName</code> in <code>hive-site.xml</code> .
Apache HCatalog	hive	hive	The WebHCat service (for REST access to Hive functionality) runs as the <code>hive</code> user.
HttpFS	httpfs	httpfs	The HttpFS service runs as this user. See HttpFS Security Configuration for instructions on how to generate the merged <code>httpfs-http.keytab</code> file.
Hue	hue	hue	Hue services run as this user.
Hue Load Balancer (Cloudera Manager 5.5 and higher)	apache	apache	The Hue Load balancer has a dependency on the <code>apache2</code> package that uses the <code>apache</code> user name. Cloudera Manager does not run processes using this user ID.
Impala	impala	impala, hive	Impala services run as this user.
Apache Kafka (CDK 1.2.0 Powered By Apache Kafka)	kafka	kafka	Kafka services run as this user.
Java KeyStore KMS (CDH 5.2.1 and higher)	kms	kms	The Java KeyStore KMS service runs as this user.
Key Trustee KMS (CDH 5.3 and higher)	kms	kms	The Key Trustee KMS service runs as this user.
Key Trustee Server (CDH 5.4 and higher)	keytrustee	keytrustee	The Key Trustee Server service runs as this user.
Kudu	kudu	kudu	Kudu services run as this user.
Llama	llama	llama	Llama runs as this user.
Apache Mahout	~	~	No special users.
MapReduce	mapred	mapred, hadoop	Without Kerberos, the JobTracker and tasks run as this user. The <code>LinuxTaskController</code> binary is owned by this user for Kerberos.
Apache Oozie	oozie	oozie	The Oozie service runs as this user.
Parquet	~	~	No special users.

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Component (Version)	Unix User ID	Groups	Functionality
Apache Pig	~	~	No special users.
Cloudera Search	solr	solr	The Solr processes run as this user.
Apache Spark	spark	spark	The Spark History Server process runs as this user.
Apache Sentry	sentry	sentry	The Sentry service runs as this user.
Apache Sqoop	sqoop	sqoop	This user is only for the Sqoop1 Metastore, a configuration option that is not recommended.
Apache Sqoop2	sqoop2	sqoop, sqoop2	The Sqoop2 service runs as this user.
Apache Whirr	~	~	No special users.
YARN	yarn	yarn, hadoop	Without Kerberos, all YARN services and applications run as this user. The LinuxContainerExecutor binary is owned by this user for Kerberos.
Apache ZooKeeper	zookeeper	zookeeper	The ZooKeeper processes run as this user. It is not configurable.

Product Compatibility Matrix for Apache Accumulo

This matrix contains compatibility information across versions of Apache Accumulo, and CDH and Cloudera Manager. For detailed information on each release, see [Apache Accumulo documentation](#).

Product	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version	Integrated into CDH
Accumulo 1.7.2	Cloudera Manager 5.0.0	CDH 5.5.0	Not Supported	Not Supported	No
Accumulo 1.6.0	Cloudera Manager 5.0.0	CDH 4.6.0 - 4.x.x, CDH 5.1.0	Not Supported	Not Supported	No
Accumulo 1.4.4	Cloudera Manager 5.0.0	CDH 4.5.0 (Not for use with CDH 5)	Not Supported	Not Supported	No
Accumulo 1.4.3	Cloudera Manager 5.0.0	CDH 4.3.0 (Not for use with CDH 5)	Not Supported	Not Supported	No

Product Compatibility Matrix for Apache Impala

This matrix contains compatibility information across versions of Impala and CDH/Cloudera Manager. For detailed information on each release, see [the Impala documentation](#).

Product	Supported Cloudera Manager Versions	Supported CDH Versions	Integrated into CDH
Impala 2.12.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.15.x	CDH 5.15.x
Impala 2.11.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.14.x	CDH 5.14.x

Product	Supported Cloudera Manager Versions	Supported CDH Versions	Integrated into CDH
Impala 2.10.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.13.x	CDH 5.13.x
Impala 2.9.x (plus additional patches) for CDH 5.12.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.12.x	CDH 5.12.x
Impala 2.8.x (plus additional patches) for CDH 5.11.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.11.x	CDH 5.11.x
Impala 2.8.x (plus a few minor patches) for CDH 5.10.x; although version string still says Impala 2.7	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.10.x	CDH 5.10.x
Impala 2.7.x for CDH 5.9.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.9.x	CDH 5.9.x
Impala 2.6.x for CDH 5.8.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.8.x	CDH 5.8.x
Impala 2.5.x for CDH 5.7.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.7.x	CDH 5.7.x
Impala 2.4.x for CDH 5.6.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.6.x	CDH 5.6.x
Impala 2.3.x for CDH 5.5.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.5.x	CDH 5.5.x
Impala 2.2.x for CDH 5.4.x	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.4.x	CDH 5.4.x
Impala 2.2.0	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.4.0	CDH 5.4.0
Impala 2.1.8	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.3.10	CDH 5.3.10
Impala 2.1.7	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.3.9	CDH 5.3.9
Impala 2.1.6	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.3.8	CDH 5.3.8
Impala 2.1.5	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.3.6	CDH 5.3.6
Impala 2.1.4	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.3.4 and CDH 5.3.5	CDH 5.3.4 and CDH 5.3.5
Impala 2.1.3	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.3.3	CDH 5.3.3
Impala 2.1.2	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.3.2	CDH 5.3.2

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Product	Supported Cloudera Manager Versions	Supported CDH Versions	Integrated into CDH
Impala 2.1.1	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.3.1	CDH 5.3.1
Impala 2.1.0	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.3.0	CDH 5.3.0
Impala 2.0.5	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.2.6	CDH 5.2.6
Impala 2.0.4	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.2.5	CDH 5.2.5
Impala 2.0.3	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.2.4	CDH 5.2.4
Impala 2.0.2	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.2.3	CDH 5.2.3
Impala 2.0.1	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.2.1	CDH 5.2.1
Impala 2.0.0	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.2.0	CDH 5.2.0
Impala 1.4.4	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.1.5	CDH 5.1.5
Impala 1.4.3	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.1.4	CDH 5.1.4
Impala 1.4.2	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.1.3	CDH 5.1.3
Impala 1.4.1	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.1.2	CDH 5.1.2
Impala 1.4.0	Cloudera Manager 5.0.0 - 5.x.x; Recommended: Cloudera Manager 5.1.0	CDH 5.1.0	CDH 5.1.0
Impala 1.3.3	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.0.5	CDH 5.0.5
Impala 1.3.2	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.0.4	CDH 5.0.4
Impala 1.3.1	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.0.1 - 5.0.x	CDH 5.0.1
Impala 1.3.0	Cloudera Manager 5.0.0 - 5.x.x	CDH 5.0.0	CDH 5.0.0

Product Compatibility Matrix for CDK Powered By Apache Kafka

CDK Powered By Apache Kafka is currently distributed as a package and in a parcel that is independent of the CDH parcel. The parcel integrates with Cloudera Manager using a Custom Service Descriptor (CSD). Using the latest features

of Kafka sometimes requires a newer version of Cloudera Manager and/or CDH that supports the feature. This matrix describes the compatibility between CDK Powered by Apache Kafka and Cloudera Manager and CDH.

For the latest documentation, see [Kafka Documentation](#).

CDK Powered By Apache Kafka Version	Upstream Kafka Version	Minimum Supported CM Version	Minimum Supported CDH Version (No Security)	Minimum Supported CDH Version (with Sentry)	Feature Notes
4.1.0	2.2.1	5.13.0	CDH 5.13.0	CDH 5.13.0	
4.0.0	2.1.0	5.13.0	CDH 5.13.0	CDH 5.13.0	
3.1.1	1.0.1	5.13.0	CDH 5.13.0	CDH 5.13.0	
3.1.0	1.0.1	5.13.0	CDH 5.13.0	CDH 5.13.0	Sentry-HA Supported
3.0.0	0.11.0	5.13.0	CDH 5.13.0	CDH 5.13.0	Sentry-HA not supported
2.2.x	0.10.2	5.9.0	CDH 5.4.0	CDH 5.9.0	
2.1.x	0.10.0	5.9.0	CDH 5.4.0	CDH 5.9.0	Sentry Authorization
2.0.x	0.9.0	5.5.3	CDH 5.4.0	CDH 5.4.0	Enhanced Security

Kafka Client-Broker Compatibility across Kafka Versions

Maintaining compatibility across different Kafka clients and brokers is a common issue. This can occur as part of any of the following scenarios:

- Upgrading your Kafka cluster without upgrading your Kafka clients
- Using a third-party application that produces to or consumes from your Kafka cluster
- Having a client program communicating with two or more Kafka clusters (e.g. consuming from one cluster and producing to a different cluster)
- Using Flume or Spark as a Kafka consumer

In these cases, it is important to understand client/broker compatibility across Kafka versions. Here are general rules that apply:

1. Newer Kafka brokers can talk to older Kafka clients. The reverse is **not** true.
2. Changes in either the major part or the minor part of the upstream version *major.minor* determines whether the client and broker are compatible.

As a result, the general pattern for upgrading Kafka from version A to version B is:

1. Change Kafka server.properties to refer to version A.
2. Upgrade the brokers to version B and restart the cluster.
3. Upgrade the clients to version B.
4. After all the clients are upgraded, change the properties from the first step to version B and restart the cluster.

Kafka Client Versions Used by Apache Flume in CDH

Flume has an embedded Kafka client which it uses to talk to Kafka clusters. Since the generally accepted practice is to have the broker running the same or newer version as the client, a CDH Flume version requires being matched to a minimum Kafka version. This matrix lists the CDK Powered By Apache Kafka (CDK) client versions used by Apache Flume

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

in CDH, as well as the underlying Apache Kafka base version. For more information about CDK, see the [Cloudera Distribution of Kafka Documentation](#).

Flume CDH Version	Minimum Supported CDK Client Version (Remote or Local)	Embedded Kafka Client Base Version
5.16.x	2.2.0	0.10.2.0
5.15.x	2.2.0	0.10.2.0
5.14.x	2.2.0	0.10.2.0
5.13.x	2.0.2	0.9.0.0
5.12.x	2.0.2	0.9.0.0
5.11.x	2.0.2	0.9.0.0
5.10.x	2.0.2	0.9.0.0
5.9.x	2.0.2	0.9.0.0
5.8.x	2.0.0	0.9.0.0
5.7.x	2.0.0	0.9.0.0

Kafka Client Versions used by Apache Spark in CDH

Spark maintains two embedded Kafka clients and can be configured to use either one. For information on compatibility with Kafka brokers, refer to the appropriate entry in the table below.

CDS Version	Embedded Kafka Client (Choose based on Kafka cluster)	Minimum Apache Kafka Broker Version	Minimum CDK Broker Version	Upstream Integration Guide	API Stability
CDS 2.4	spark-streaming-kafka-0.10	0.10.0	Kafka 2.1	Spark 2.4 + Kafka 0.10	Experimental
	spark-streaming-kafka-0.8	0.8.2.1	Kafka 2.0	Spark 2.4 + Kafka 0.8	Stable
CDS 2.3	spark-streaming-kafka-0.10	0.10.0	Kafka 2.1	Spark 2.3 + Kafka 0.10	Experimental
	spark-streaming-kafka-0.8	0.8.2.1	Kafka 2.0	Spark 2.3 + Kafka 0.8	Stable
CDS 2.2	spark-streaming-kafka-0.10	0.10.0	Kafka 2.1	Spark 2.2 + Kafka 0.10	Experimental
	spark-streaming-kafka-0.8	0.8.2.1	Kafka 2.0	Spark 2.2 + Kafka 0.8	Stable
CDS 2.1	spark-streaming-kafka-0.10	0.10.0	Kafka 2.1	Spark 2.1 + Kafka 0.10	Experimental
	spark-streaming-kafka-0.8	0.8.2.1	Kafka 2.0	Spark 2.1 + Kafka 0.8	Stable
CDS 2.0	spark-streaming-kafka-0.8	0.8.2.x	Kafka 2.0	Spark 2.0 + Kafka	Stable

Product Compatibility Matrix for Apache Kudu

This matrix contains compatibility information across versions of Kudu and CDH/Cloudera Manager. Starting with Apache Kudu 1.5.0 / CDH 5.13.0, Kudu is fully integrated into CDH and ships as part of the CDH parcel and packages.

Documentation for Kudu 1.5.0 / CDH 5.13.x (and higher) can be found at [Apache Kudu Guide](#). Documentation for all lower versions can be found at [Apache Kudu 1.4.0 \(and lower\)](#).

Product	Feature	Lowest Supported Cloudera Manager Version	Supported CDH Versions	Integrated into CDH
Kudu 1.7.0 / CDH 5.15.x	Fully integrated into the CDH 5.15 parcel and packages.	Cloudera Manager 5.15.0	CDH 5.15.x	Yes
Kudu 1.6.0 / CDH 5.14.x	Fully integrated into the CDH 5.14 parcel and packages.	Cloudera Manager 5.14.0	CDH 5.14.x	Yes
Kudu 1.5.0 / CDH 5.13.x	Fully integrated into the CDH 5.13 parcel and packages.	Cloudera Manager 5.13.0	CDH 5.13.x	Yes
Kudu 1.4.0 / CDH 5.12.x	Available as a CSD within Cloudera Manager.	Cloudera Manager 5.12.0	CDH 5.12.x	No
Kudu 1.3.0 / CDH 5.11.x	Available as a CSD within Cloudera Manager.	Cloudera Manager 5.11.0	CDH 5.11.x	No
Kudu 1.2.0 / CDH 5.10.x	Distributed both as package and parcel. Kudu's first generally available (GA) release that is supported on Cloudera Manager and CDH 5.10.x clusters.	Cloudera Manager 5.10.0	CDH 5.10.x	No
Kudu 1.1.x (and lower) - Beta	Public beta releases that are not commercially supported on CDH clusters.	N/A	N/A	No

Supported Operating Systems (Kudu 1.2.0 / CDH 5.10.x and higher)

Starting in CDH 5.13, Kudu supports the same OSs and versions as Cloudera Manager and CDH with the following exceptions.

- RHEL 5 is not supported.
- Ubuntu 12 is not supported.
- Debian 7 is not supported.
- macOS
 - 10.10 Yosemite, 10.11 El Capitan, and 10.12 Sierra are supported.
 - Pre-built macOS packages are not provided.
- Microsoft Windows is not supported.

For a complete list of supported versions, see [CDH and Cloudera Manager Supported Operating Systems](#) on page 20.

Filesystem Requirements

Kudu is supported on ext4 and XFS. Kudu requires a kernel version and filesystem that supports hole punching. Hole punching is the use of the `fallocate(2)` system call with the `FALLOC_FL_PUNCH_HOLE` option set. For more details, see Kudu troubleshooting - [Error During Hole Punch Test](#).

nscd Recommendation

Although not a strict requirement, it's highly recommended that you use `nscd` to cache both DNS name resolution and static name resolution for Kudu. See [Troubleshooting Kudu](#) for more information.

JDK Requirements

- JDK 7
- JDK 8

Note that Spark 2.2 (and higher) requires Java 8 at runtime even though the Kudu Spark 2.x integration is Java 7 compatible. Spark 2.2 is the default dependency version as of Kudu 1.5.0 / CDH 5.13.

Hardware Requirements

- Kudu currently requires a CPU that supports the SSSE3 and SSE4.2 instruction sets. If you are to run Kudu inside a VM, enable SSE4.2 pass-through to pass through SSE4.2 support into the VM.
- One or more hosts to run Kudu masters. You should have either one master (provides no fault tolerance), three masters (can tolerate one failure), or five masters (can tolerate two failures).
- One or more hosts to run Kudu tablet servers. With replication, a minimum of three tablet servers is necessary.
- **Storage** - If solid state storage is available, storing Kudu WALs on such high-performance media might significantly improve latency when Kudu is configured for its highest durability levels.

Product Compatibility Matrix for Cloudera Navigator

This matrix contains compatibility information across versions of Cloudera Navigator, Cloudera Manager, and CDH. For detailed information on each release, see [Cloudera Navigator documentation](#).



Note:

Cloudera Navigator requires HiveServer2 for complete governance Hive queries. Cloudera Navigator does not capture audit events for queries that are run on HiveServer1/Hive CLI, and lineage is not captured for certain types of operations that are run on HiveServer1.

[HiveServer1 and Hive CLI are deprecated in CDH 5](#). If you use Cloudera Navigator to capture auditing, lineage, and metadata for Hive operations, upgrade to HiveServer2 if you have not done so already.

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
Cloudera Navigator 2.15.x	Auditing, Metadata, Analytics, and Security	5.16.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> – HDFS, HBase - 5.0.0 – Hue - 5.5.0 for LDAP 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<ul style="list-style-type: none"> user operations – HiveServer2 - 5.0.0 – Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> – HDFS, HiveServer2, MapReduce, Oozie, Sqoop 1 - 5.0.0 – Pig - 5.0.0 – YARN - 5.0.0 – Impala - 5.4.0 – Spark - 1.6, 2.3 		
Cloudera Navigator 2.14.x	Auditing, Metadata, Analytics, and Security	5.15.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> – HDFS, HBase - 5.0.0 – Hue - 5.5.0 for LDAP user operations – HiveServer2 - 5.0.0 – Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> – HDFS, HiveServer2, MapReduce, Oozie, Sqoop 1 - 5.0.0 – Pig - 5.0.0 – YARN - 5.0.0 – Impala - 5.4.0 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<ul style="list-style-type: none"> - Spark - 1.6, 2.3 		
Cloudera Navigator 2.13.x	Auditing, Metadata, Analytics, and Security	5.14.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 5.0.0 - Hue - 5.5.0 for LDAP user operations - HiveServer2 - 5.0.0 - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, HiveServer2, MapReduce, Oozie, Sqoop 1 - 5.0.0 - Pig - 5.0.0 - YARN - 5.0.0 - Impala - 5.4.0 - Spark - 1.6, 2.3 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0
Cloudera Navigator 2.12.x	Auditing, Metadata, Analytics, and Security	5.13.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.4.0, 5.5.0 for LDAP user operations - HiveServer2 - 4.2.0, 4.4.0 for operations denied due to lack of privileges. 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<ul style="list-style-type: none"> - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, HiveServer2, MapReduce, Oozie, Sqoop 1 - 4.4.0 - Pig - 4.6.0 - YARN - 5.0.0 - Impala - 5.4.0 - Spark - 1.6 		
Cloudera Navigator 2.11.x	Auditing, Metadata, Analytics, and Security	5.12.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.4.0, 5.5.0 for LDAP user operations - HiveServer2 - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, HiveServer2, MapReduce, Oozie, Sqoop 1 - 4.4.0 - Pig - 4.6.0 - YARN - 5.0.0 - Impala - 5.4.0 - Spark - 1.6 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
Cloudera Navigator 2.10.x	Auditing, Metadata, Analytics, and Security	5.11.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> – HDFS, HBase - 4.0.0 – Hue - 4.4.0, 5.5.0 for LDAP user operations – HiveServer2 - 4.2.0, 4.4.0 for operations denied due to lack of privileges. – Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> – HDFS, HiveServer2, MapReduce, Oozie, Sqoop 1 - 4.4.0 – Pig - 4.6.0 – YARN - 5.0.0 – Impala - 5.4.0 – Spark - 1.6 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0
Cloudera Navigator 2.9.x	Auditing, Metadata, Analytics, and Security	5.10.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> – HDFS, HBase - 4.0.0 – Hue - 4.4.0, 5.5.0 for LDAP user operations – HiveServer2 - 4.2.0, 4.4.0 for operations denied due to lack of privileges. 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<ul style="list-style-type: none"> - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, HiveServer2, MapReduce, Oozie, Sqoop 1 - 4.4.0 - Pig - 4.6.0 - YARN - 5.0.0 - Impala - 5.4.0 - Spark (available but unsupported) - 5.4.0 		
Cloudera Navigator 2.8.x	Auditing, Metadata, Analytics, and Security	5.9.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.4.0, 5.5.0 for LDAP user operations - HiveServer2 - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, HiveServer2, MapReduce, Oozie, Sqoop 1 - 4.4.0 - Pig - 4.6.0 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<ul style="list-style-type: none"> – YARN - 5.0.0 – Impala - 5.4.0 – Spark (available but unsupported) - 5.4.0 		
Cloudera Navigator 2.7.x	Auditing, Metadata, Analytics, and Security	5.8.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> – HDFS, HBase - 4.0.0 – Hue - 4.4.0, 5.5.0 for LDAP user operations – Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. – Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> – HDFS, Hive, MapReduce, Oozie, Sqoop 1 - 4.4.0 – Pig - 4.6.0 – YARN - 5.0.0 – Impala - 5.4.0 – Spark (available but unsupported)- 5.4.0 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0
Cloudera Navigator 2.6.x	Auditing, Metadata, Analytics, and Security	5.7.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> – HDFS, HBase - 4.0.0 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<ul style="list-style-type: none"> - Hue - 4.4.0, 5.5.0 for LDAP user operations - Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, Hive, MapReduce, Oozie, Sqoop 1 - 4.4.0 - Pig - 4.6.0 - YARN - 5.0.0 - Impala - 5.4.0 - Spark (available but unsupported)- 5.4.0 		
Cloudera Navigator 2.5.x	Auditing, Metadata, Analytics, and Security	5.6.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.2.0 - Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, Hive, Impala, 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<p>MapReduce, Oozie, Sqoop 1 - 4.4.0</p> <ul style="list-style-type: none"> - Pig - 4.6.0 - YARN - 5.0.0 - Impala and Spark (available but unsupported)- 5.4.0 		
Cloudera Navigator 2.4.x	Auditing, Metadata, Analytics, and Security	5.5.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.2.0 - Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, Hive, Impala, MapReduce, Oozie, Sqoop 1 - 4.4.0 - Pig - 4.6.0 - YARN - 5.0.0 - Impala and Spark (available but unsupported)- 5.4.0 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0
Cloudera Navigator 2.3.x	Auditing, Metadata, and Security	5.4.0	<ul style="list-style-type: none"> • Audit Component 	Impala 1.3.0 with CDH 5.0.0	CDH 5.4.0

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.2.0 - Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, Hive, Impala, MapReduce, Oozie, Sqoop 1 - 4.4.0 - Pig - 4.6.0 - YARN - 5.0.0 - Impala and Spark (available but unsupported)- 5.4.0 		
Cloudera Navigator 2.2.x	Auditing, Metadata, and Security	5.3.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.2.0 - Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, Hive, Oozie, 	Impala 1.3.0 with CDH 5.0.0	Not Supported

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<p>MapReduce, Sqoop 1 - 4.4.0</p> <ul style="list-style-type: none"> - Pig - 4.6.0 - YARN - 5.0.0 		
Cloudera Navigator 2.1.x	Auditing, Metadata, and Security	5.2.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.2.0 - Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, Hive, Oozie, MapReduce, Sqoop 1 - 4.4.0 - Pig - 4.6.0 - YARN - 5.0.0 	Impala 1.3.0 with CDH 5.0.0	Not Supported
Cloudera Navigator 2.0.1	Auditing, Metadata, and Security	5.1.2	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.2.0 - Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component 	Impala 1.3.0 with CDH 5.0.0	Not Supported

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version
			<ul style="list-style-type: none"> - HDFS, Hive, Oozie, MapReduce, Sqoop 1 - 4.4.0 - Pig - 4.6.0 - YARN - 5.0.0 		
Cloudera Navigator 2.0.0	Auditing, Metadata, and Security	5.1.0	<ul style="list-style-type: none"> • Audit Component <ul style="list-style-type: none"> - HDFS, HBase - 4.0.0 - Hue - 4.2.0 - Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. - Sentry - 5.1.0 • Metadata Component <ul style="list-style-type: none"> - HDFS, Hive, Oozie, MapReduce, Sqoop 1 - 4.4.0 - Pig - 4.6.0 - YARN - 5.0.0 	Impala 1.2.1 with CDH 4.4.0	Not Supported
Cloudera Navigator 1.2.x	Auditing	5.0.0	<ul style="list-style-type: none"> • HDFS, HBase - 4.0.0 • Hue - 4.2.0 • Hive - 4.2.0, 4.4.0 for operations denied due to lack of privileges. 	Impala 1.3.0 with CDH 5.0.0	Not Supported
	Metadata (2.0 beta 2)	5.0.0	<ul style="list-style-type: none"> • HDFS, Hive, Oozie, MapReduce, Sqoop 1 - 4.4.0 • Pig - 4.6.0 	Not Supported	Not Supported

Cloudera Navigator Supported Databases

Cloudera Navigator supports the same databases as CDH and Cloudera Manager. For more information, see [CDH and Cloudera Manager Supported Databases](#) on page 31.

Cloudera Navigator Supported Browsers

The Cloudera Navigator console supports the following browsers:

- Mozilla Firefox 24 and higher
- Google Chrome 36 and higher
- Internet Explorer 11
- Safari 5 and higher

Cloudera Navigator Supported CDH and Managed Service Versions

This section describes the CDH and managed service versions supported by the Cloudera Navigator auditing and metadata features.

Cloudera Navigator Auditing

This section describes the audited operations and service versions supported by Cloudera Navigator auditing.

Component	Operations (For details, see Service Audit Events).	Minimum Supported Service Version
HDFS	<ul style="list-style-type: none"> • Operations that access or modify a file's or directory's data or metadata • Operations denied due to lack of privileges 	CDH 5.0.0
HBase	<ul style="list-style-type: none"> • In simple authentication mode, if the HBase Secure RPC Engine property is <code>false</code> (the default), the username in log events is <code>UNKNOWN</code>. To see a meaningful username: <ol style="list-style-type: none"> 1. Click the HBase service. 2. Click the Configuration tab. 3. Select Service-wide > Security. 4. Set the HBase Secure RPC Engine property to <code>true</code>. 5. Save the change and restart the service. 	CDH 5.0.0
Hive	<ul style="list-style-type: none"> • Operations (except grant, revoke, and metadata access only) sent to HiveServer2 • Operations denied due to lack of privileges <p>Limitations:</p> <ul style="list-style-type: none"> • Actions taken against Hive using the Hive CLI are <i>not</i> audited. Therefore if you have enabled auditing you should disable the Hive CLI to prevent actions against Hive that are not audited. • In simple authentication mode, the username in log events is the username passed in the HiveServer2 connect command. If you do not pass a username in the connect command, the username in log events is <code>anonymous</code>. 	CDH 5.0.0
Hue	<ul style="list-style-type: none"> • Operations (except grant, revoke, and metadata access only) sent through the Beeswax Server • User operations such as log in, log out, add and remove user, add and remove LDAP group, add and remove user from LDAP group 	CDH 5.0.0

Component	Operations (For details, see Service Audit Events).	Minimum Supported Service Version
Impala	<ul style="list-style-type: none"> Queries denied due to lack of privileges Queries that pass analysis 	Impala 1.3.0 with CDH 5.0.0
Navigator Metadata Server	<ul style="list-style-type: none"> Viewing and changing audit reports Viewing and changing authorization configurations Viewing and changing metadata Viewing and changing policies Viewing and changing saved searches 	Cloudera Navigator 2.3
Sentry	<ul style="list-style-type: none"> Operations sent to the HiveServer2 and Hive Metastore Server roles and Impala service Adding and deleting roles, assigning roles to groups and removing roles from groups, creating and deleting privileges, granting and revoking privileges Operations denied due to lack of privileges <p>You do not directly configure the Sentry service for auditing. Instead, when you configure the Hive and Impala services for auditing, grant, revoke, and metadata operations appear in the Hive or Impala service audit logs.</p>	CDH 5.1.0
Solr	<ul style="list-style-type: none"> Index creation and deletion Schema and configuration file modification Index, service, document tag access 	CDH 5.4.0

Cloudera Navigator Metadata

This section describes the CDH and managed service versions supported by the Cloudera Navigator metadata feature.

Component	Minimum Supported Version
HDFS. However, federated HDFS is <i>not supported</i> .	CDH 5.0.0
Hive	CDH 5.0.0
Impala	CDH 5.4.0
MapReduce	CDH 5.0.0
Oozie. Supported actions: <ul style="list-style-type: none"> 2.4 - map-reduce, pig, hive, hive2, sqoop 2.3 and lower - map-reduce, pig, hive, sqoop Unsupported actions include email, shell, and ssh.	CDH 5.0.0
Pig	CDH 5.0.0
Spark	CDH 5.4.0
Sqoop 1. All Cloudera connectors are supported.	CDH 5.0.0
YARN	CDH 5.0.0

Data at Rest Encryption Requirements

Encryption comprises several components, each with its own requirements.

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Data at rest encryption protection can be applied at a number of levels within Hadoop:

- OS filesystem-level
- Network-level
- HDFS-level (protects both data at rest and in transit)

This section contains the various hardware and software requirements for all encryption products used for Data at Rest Encryption.

For more information on supported operating systems, see [Product Compatibility Matrix for Cloudera Navigator Encryption](#) on page 63.

For more information on the components, concepts, and architecture for encrypting data at rest, see [Encrypting Data at Rest](#).

Entropy Requirements

See [Entropy Requirements](#).

Cloudera Manager Requirements

Installing and managing Key Trustee Server using Cloudera Manager requires Cloudera Manager 5.4.0 and higher. Key Trustee Server does not require Cloudera Navigator Audit Server or Metadata Server.

umask Requirements

Key Trustee Server installation requires the default `umask` of `0022`.

Network Requirements

For new Key Trustee Server installations (5.4.0 and higher) and migrated upgrades (see [Migrate Apache Web Server to CherryPy](#) for more information), Key Trustee Server requires the following TCP ports to be opened for inbound traffic:

- 11371
Clients connect to this port over HTTPS.
- 11381 (PostgreSQL)
The passive Key Trustee Server connects to this port for database replication.

For upgrades that are not migrated to the CherryPy web server, the pre-upgrade port settings are preserved:

- 80
Clients connect to this port over HTTP to obtain the Key Trustee Server public key.
- 443 (HTTPS)
Clients connect to this port over HTTPS.
- 5432 (PostgreSQL)
The passive Key Trustee Server connects to this port for database replication.

TLS Certificate Requirements

To ensure secure network traffic, Cloudera recommends obtaining Transport Layer Security (TLS) certificates specific to the hostname of your Key Trustee Server. To obtain the certificate, generate a Certificate Signing Request (CSR) for the fully qualified domain name (FQDN) of the Key Trustee Server host. The CSR must be signed by a trusted Certificate Authority (CA). After the certificate has been verified and signed by the CA, the Key Trustee Server TLS configuration requires:

- The CA-signed certificate
- The private key used to generate the original CSR
- The intermediate certificate/chain file (provided by the CA)

Cloudera recommends not using self-signed certificates. If you use self-signed certificates, you must use the `--skip-ssl-check` parameter when registering Navigator Encrypt with the Key Trustee Server. This skips TLS hostname validation, which safeguards against certain network-level attacks. For more information regarding insecure mode, see [Table 1](#).

Product Compatibility Matrix for Cloudera Navigator Encryption

Cloudera Navigator encryption comprises several components.

Although the version numbers differ between some Cloudera Navigator encryption components and Cloudera Enterprise, the encryption component releases are generally coordinated with Cloudera Enterprise releases. See the following table for information on which component versions correspond with a given Cloudera Enterprise release.



Note: The following table is not a compatibility matrix. It lists the Cloudera Navigator encryption component versions that were released around the same time as the corresponding Cloudera Enterprise release.

Table 11: Cloudera Navigator Encryption Components Release Versions

Cloudera Enterprise Release	Cloudera Navigator Key Trustee Server Version	Key Trustee KMS Version	Cloudera Navigator Key HSM Version	Cloudera Navigator Encrypt Version
5.16.x	<i>No release; use 5.15.x</i>	5.16.x	<i>No release; use 1.11.x</i>	3.16.x
5.15.x	5.15.x	5.15.x	<i>No release; use 1.11.x</i>	3.15.x
5.14.x	5.14.x	5.14.x	1.11.x	<i>No release; use 3.14.x</i>
5.13.x	5.13.x	5.13.x	<i>No release; use 1.10.x</i>	3.14.x
5.12.x	5.12.x	5.12.x	1.10.x	3.13.x
5.11.x	5.11.x	5.11.x	1.9.x	3.12.x
5.10.x	5.10.x	5.10.x	<i>No release; use 1.8.x</i>	3.11.x
5.9.x	5.9.x	5.9.x	1.8.x	<i>No release; use 3.10.x</i>
5.8.x	5.8.x	5.8.x	1.7.x	3.10.x
5.7.x	5.7.x	5.7.x	1.6.x	3.9.x
5.6.x	<i>No release; use 5.5.x</i>	<i>No release; use 5.5.x</i>	<i>No release; use 1.5.x</i>	<i>No release; use 3.8.x</i>
5.5.x	5.5.x	5.5.x	1.5.x	3.8.x
5.4.x	5.4.x	5.4.x	1.4.x	3.7.x
5.3.x	3.8.x	5.3.x	1.3.x	3.6.x

See below for the individual compatibility matrices for each component:

Cloudera Navigator Key Trustee Server

Because of a change in the ports used by Key Trustee Server, Navigator Encrypt versions lower than 3.7 and Key Trustee KMS versions lower than 5.4 are not supported in Key Trustee Server 5.4 and higher.

Key Trustee Server: Recommended Hardware and Supported Distributions

Recommended Hardware and Supported Distributions

Key Trustee Server must be installed on a dedicated server or virtual machine (VM) that is not used for any other purpose. The backing PostgreSQL database must be installed on the same host as the Key Trustee Server, and must not be shared with any other services. For high availability, the active and passive Key Trustee Servers must not share physical resources. See [Resource Planning for Data at Rest Encryption](#) for more information.

The recommended minimum hardware specifications are as follows:

- Processor: 1 GHz 64-bit quad core
- Memory: 8 GB RAM
- Storage: 20 GB on moderate-to high-performance disk drives

Table 12: Cloudera Navigator Key Trustee Server Compatibility Matrix

Cloudera Navigator Key Trustee Server Version	Supported Operating Systems	Supported Cloudera Manager Versions	Supported Cloudera Navigator Key HSM Versions	Supported Key Trustee KMS Versions	Supported Cloudera Navigator Encrypt Versions
5.15.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.10, 7.1 - 7.6 • Oracle Linux: 6.7 - 6.9, 7.1 - 7.5 	5.15.x and higher	1.4.x and higher	5.4.x and higher	3.7.x and higher
5.14.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.10, 7.1 - 7.6 • Oracle Linux: 6.7 - 6.9, 7.1 - 7.5 	5.14.x and higher	1.4.x and higher	5.4.x and higher	3.7.x and higher
5.13.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.10, 7.1 - 7.6 • Oracle Linux: 6.7 - 6.9, 7.1 - 7.5 	5.13.x and higher	1.4.x and higher	5.4.x and higher	3.7.x and higher
5.12.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.10, 7.1 - 7.5 • Oracle Linux: 6.7 - 6.9, 7.1 - 7.5 	5.12.x and higher	1.4.x and higher	5.4.x and higher	3.7.x and higher
5.11.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.8, 7.1 - 7.5 • Oracle Linux: 6.7, 6.8, 7.1 - 7.5 	5.11.x and higher	1.3.x and higher	5.4.x and higher	3.7.x and higher

Cloudera Navigator Key Trustee Server Version	Supported Operating Systems	Supported Cloudera Manager Versions	Supported Cloudera Navigator Key HSM Versions	Supported Key Trustee KMS Versions	Supported Cloudera Navigator Encrypt Versions
5.10.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.8, 7.1 - 7.3 • Oracle Linux: 6.7, 6.8, 7.1 - 7.3 	5.10.x and higher	1.3.x and higher	5.4.x and higher	3.7.x and higher
5.9.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.8, 7.1, 7.2 • Oracle Linux: 6.7, 6.8, 7.1, 7.2 	5.9.x	1.3.x and higher	5.4.x and higher	3.7.x and higher
5.8.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.7, 7.1, 7.2 • Oracle Linux: 6.7, 7.1, 7.2 	5.8.x, 5.9.x	1.3.x and higher	5.4.x and higher	3.7.x and higher
5.7.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.7, 7.1, 7.2 • Oracle Linux: 6.7, 7.1, 7.2 	5.7.x - 5.9.x	1.3.x and higher	5.4.x and higher	3.7.x and higher
5.5.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.7, 7.1 • Oracle Linux: 6.7, 7.1, 7.2 	5.5.x - 5.9.x	1.3.x and higher	5.4.x and higher	3.7.x and higher
5.4.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4 - 6.6 	5.4.x - 5.9.x	1.3.x and higher	5.4.x and higher	3.7.x and higher
3.8.x	<ul style="list-style-type: none"> • RHEL and CentOS: 6.4, 6.5 	Not supported	1.3.x	5.3.x	3.6.x

Key Trustee KMS

Key Trustee KMS: Recommended Hardware and Supported Distributions

The recommended minimum hardware specifications are as follows:

- Oracle Java Runtime Environment (JRE) 7 or higher with Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files:
 - [JCE for Java SE 7](#)
 - [JCE for Java SE 8](#)



Note: JDK 1.8u161 and higher enable unlimited strength encryption by default, and do not require policy files.

- Processor: 2 GHz 64-bit quad core
- Memory: 16 GB RAM
- Storage: 40 GB on moderate-to high-performance disk drives

The Key Trustee KMS workload is CPU-intensive. Cloudera recommends using machines with capabilities equivalent to your NameNode hosts, with Intel CPUs that support [AES-NI](#) for optimum performance. Also, Cloudera strongly recommends that you enable TLS for both the HDFS and the Key Trustee KMS services to prevent the passage of plain text key material between the KMS and HDFS data nodes.

Table 13: Key Trustee KMS Compatibility Matrix

Key Trustee KMS Version	Supported Operating Systems	Supported Key Trustee Server Versions	Supported Cloudera Manager Versions	Supported CDH Versions
5.16.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.10, 7.1 - 7.6 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.9, 7.1 - 7.5 • SLES: 11 SP2 - SP4, 12 SP1 - SP3 • Debian: 7.1, 7.8, 8.2, 8.4 • Ubuntu: 14.04, 16.04 	5.4.x and higher	5.16.x and higher	5.3.x - 5.16.x
5.15.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.10, 7.1 - 7.6 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.9, 7.1 - 7.5 • SLES: 11 SP2 - SP4, 12 SP1 - SP3 • Debian: 7.1, 7.8, 8.2, 8.4 • Ubuntu: 12.04, 14.04, 16.04 	5.4.x and higher	5.15.x and higher	5.3.x - 5.15.x
5.14.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.10, 7.1 - 7.6 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.9, 7.1 - 7.5 	5.4.x and higher	5.14.x and higher	5.3.x - 5.14.x

Key Trustee KMS Version	Supported Operating Systems	Supported Key Trustee Server Versions	Supported Cloudera Manager Versions	Supported CDH Versions
	<ul style="list-style-type: none"> • SLES: 11 SP2 - SP4, 12 SP1, 12 SP2 • Debian: 7.1, 7.8, 8.2, 8.4 • Ubuntu: 12.04, 14.04, 16.04 			
5.13.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.10, 7.1 - 7.6 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.9, 7.1 - 7.5 • SLES: 11 SP2 - SP4, 12 SP1, 12 SP2 • Debian: 7.1, 7.8, 8.2, 8.4 • Ubuntu: 12.04, 14.04, 16.04 	5.4.x and higher	5.13.x and higher	5.3.x - 5.13.x
5.12.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.10, 7.1 - 7.5 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.9, 7.1 - 7.5 • SLES: 11 SP2 - SP4, 12 SP1, 12 SP2 • Debian: 7.1, 7.8, 8.2, 8.4 • Ubuntu: 12.04, 14.04, 16.04 	5.4.x and higher	5.12.x and higher	5.3.x - 5.12.x
5.11.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.8, 7.1 - 7.5 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.8, 7.1 - 7.5 • SLES: 11 SP2 - SP4, 12 SP1, 12 SP2 • Debian: 7.1, 7.8, 8.2, 8.4 	5.4.x and higher	5.11.x and higher	5.3.x - 5.11.x

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Key Trustee KMS Version	Supported Operating Systems	Supported Key Trustee Server Versions	Supported Cloudera Manager Versions	Supported CDH Versions
	<ul style="list-style-type: none"> • Ubuntu: 12.04, 14.04, 16.04 			
5.10.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.8, 7.1 - 7.3 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.8, 7.1 - 7.3 • SLES: 11 SP2 - SP4, 12 SP1 • Debian: 7.1, 7.8, 8.2, 8.4 • Ubuntu: 12.04, 14.04 	5.4.x and higher	5.10.x and higher	5.3.x - 5.10.x
5.9.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.8, 7.1, 7.2 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.8, 7.1, 7.2 • SLES: 11 SP2 - SP4, 12 SP1 • Debian: 7.1, 7.8, 8.2, 8.4 • Ubuntu: 12.04, 14.04 	5.4.x and higher	5.9.x and higher	5.3.x - 5.9.x
5.8.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.7, 7.1, 7.2 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.7, 7.1, 7.2 • SLES: 11 SP2 - SP4 • Debian: 7.1, 7.8, 8.2 • Ubuntu: 12.04, 14.04 	5.4.x and higher	5.8.x and higher	5.3.x - 5.8.x
5.7.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.7, 7.1, 7.2 • Oracle Linux: 5.7, 5.10, 5.11, 6.4 - 6.7, 7.1, 7.2 	5.4.x and higher	5.7.x and higher	5.3.x - 5.7.x

Key Trustee KMS Version	Supported Operating Systems	Supported Key Trustee Server Versions	Supported Cloudera Manager Versions	Supported CDH Versions
	<ul style="list-style-type: none"> • SLES: 11 SP2 - SP4 • Debian: 7.1, 7.8 • Ubuntu: 12.04, 14.04 			
5.5.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.7, 7.1 • Oracle Linux: 5.7, 5.10, 6.4 - 6.7, 7.1 • SLES: 11 SP2, SP3 • Debian: 7.1 • Ubuntu: 12.04, 14.04 	5.4.x and higher	5.5.x and higher	5.3.x - 5.5.x
5.4.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.6 • Oracle Linux: 5.7, 5.10, 6.4 - 6.6 • SLES: 11 SP2, SP3 • Debian: 7.1 • Ubuntu: 12.04, 14.04 	5.4.x and higher	5.4.x and higher	5.3.x, 5.4.x
5.3.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 6.4, 6.5 • Oracle Linux: 5.7, 6.4, 6.5 • SLES: 11 SP2, 11 SP3 • Debian: 7.1 • Ubuntu: 12.04, 14.04 	3.8.x	5.3.x and higher	5.3.x

Cloudera Navigator HSM KMS

Navigator HSM KMS: Recommended Hardware and Supported Distributions

The recommended minimum hardware specifications are as follows:

- Oracle Java Runtime Environment (JRE) 7 or higher with Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files:
 - [JCE for Java SE 7](#)
 - [JCE for Java SE 8](#)



Note: JDK 1.8u161 and higher enable unlimited strength encryption by default, and do not require policy files.

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

- Processor: 2 GHz 64-bit quad core
- Memory: 16 GB RAM
- Storage: 40 GB on moderate-to high-performance disk drives

Supported HSM devices:

- SafeNet Luna
 - HSM software version: 6.2.2-5
 - HSM firmware version: 6.10.9
 - Client: 6.2.2
- Thales nShield Solo, nShield Connect
 - Server version: 3.67.11cam4
 - Firmware: 2.65.2
 - Security World Version: 12.30

Table 14: HSM KMS Compatibility Matrix

HSM KMS Version	Supported Operating Systems	Supported Cloudera Manager Version	Supported CDH Versions
5.16.x	<ul style="list-style-type: none">• RHEL and CentOS: 6.8, 6.9, 7.1 - 7.5	5.16.x and higher	5.12.x - 5.16.x
5.15.x	<ul style="list-style-type: none">• RHEL and CentOS: 6.8, 6.9, 7.1 - 7.5	5.15.x and higher	5.12.x - 5.15.x
5.14.x	<ul style="list-style-type: none">• RHEL and CentOS: 6.8, 6.9, 7.1 - 7.4	5.14.x and higher	5.12.x - 5.14.x
5.13.x	<ul style="list-style-type: none">• RHEL and CentOS: 6.8, 6.9, 7.1 - 7.4	5.13.x and higher	5.12.x, 5.13.x
5.12.x	<ul style="list-style-type: none">• RHEL and CentOS: 6.8, 6.9, 7.1 - 7.3	5.12.x and higher	5.12.x

Cloudera Navigator Key HSM

Cloudera Navigator Key HSM must be installed on the same host as Key Trustee Server. Although Key HSM is compatible across all versions of Key Trustee Server, Cloudera strongly recommends also upgrading Key HSM after you upgrade Key Trustee Server.

Key HSM: Recommended Hardware and Supported Distributions

The following are prerequisites for installing Navigator Key HSM:


- Oracle Java Runtime Environment (JRE) 7 or higher with Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files:
 - [JCE for Java SE 7](#)
 - [JCE for Java SE 8](#)



Note: JDK 1.8u161 and higher enable unlimited strength encryption by default, and do not require policy files.

- A supported HSM device:

- Thales (formerly SafeNet) Luna
 - v6
 - HSM firmware version: 6.2.1
 - HSM software version: 5.2.3-1
 - v7
 - HSM firmware version: 7.0.3
 - HSM software version: 7.2.0
 - HSM firmware version: 6.2.1
 - HSM software version: 5.2.3-1
 - SafeNet KeySecure
 - HSM firmware version: 6.2.1
 - HSM software version: 8.0.1, 8.1.0, 8.7.0
 - Thales nShield Solo, nShield Connect
 - HSM firmware version: 11.4.0
 - Client software version: 2.28.9cam136
- Key Trustee Server 3.8 or higher

 **Important:** You must install Key HSM on the same host as Key Trustee Server.

Root access is required to install Navigator Key HSM.

Table 15: Cloudera Navigator Key HSM Compatibility Matrix

Cloudera Navigator Key HSM Version	Supported Operating Systems	Supported Key Trustee Server Versions
1.11.x	• RHEL and CentOS: 6.4 - 6.9, 7.1 - 7.6	3.8.x and higher
1.10.x	• RHEL and CentOS: 6.4 - 6.9, 7.1 - 7.3	3.8.x and higher
1.9.x	• RHEL and CentOS: 6.4 - 6.8, 7.1, 7.2	3.8.x and higher
1.8.x	• RHEL and CentOS: 6.4 - 6.8, 7.1, 7.2	3.8.x and higher
1.7.x	• RHEL and CentOS: 6.4 - 6.7, 7.1, 7.2	3.8.x and higher
1.6.x	• RHEL and CentOS: 6.4 - 6.7, 7.1, 7.2	3.8.x and higher
1.5.x	• RHEL and CentOS: 6.4 - 6.7, 7.1	3.8.x and higher
1.4.x	• RHEL and CentOS: 6.4 - 6.6	3.8.x and higher
1.3.x	• RHEL and CentOS: 6.4, 6.5	3.8.x and higher

Cloudera Navigator Encrypt



Note: Cloudera Enterprise, with the exception of Cloudera Navigator Encrypt, is supported on platforms with Security-Enhanced Linux (SELinux) enabled and in `enforcing` mode. Cloudera is not responsible for SELinux policy development, support, or enforcement. If you experience issues running Cloudera software with SELinux enabled, contact your OS provider for assistance.

If you are using SELinux in `enforcing` mode, Cloudera Support can request that you disable SELinux or change the mode to `permissive` to rule out SELinux as a factor when investigating reported issues.

Supported command-line interpreters:

- `sh` (Bourne)
- `bash` (Bash)
- `dash` (Debian)



Note: Navigator Encrypt does not support installation or use in `chroot` environments.

Network Requirements

For new Navigator Key Trustee Server (5.4.0 and higher) installations, Navigator Encrypt initiates TCP traffic over port 11371 (HTTPS) to the Key Trustee Server.

For upgrades and Key Trustee Server versions lower than 5.4.0, Navigator Encrypt initiates TCP traffic over ports 80 (HTTP) and 443 (HTTPS) to the Navigator Key Trustee Server.

Internet Access

You must have an active connection to the Internet to download many package dependencies, unless you have internal repositories or mirrors containing the dependent packages.

Maintenance Window

Data is not accessible during the encryption process. Plan for system downtime during installation and configuration.

Administrative Access

To enforce a high level of security, all Navigator Encrypt commands require administrative (root) access (including installation and configuration). If you do not have administrative privileges on your server, contact your system administrator before proceeding.

Package Dependencies

Navigator Encrypt requires these packages, which are resolved by your distribution package manager during installation:

- `dkms`
- `keyutils`
- `ecryptfs-utils`
- `libkeytrustee`
- `navencrypt-kernel-module`
- `openssl`
- `lsuf`
- `gcc`
- `cryptsetup`

These packages may have other dependencies that are also resolved by your package manager. Installation works with gcc, gcc3, and gcc4.

Table 16: Cloudera Navigator Encrypt Compatibility Matrix

Cloudera Navigator Encrypt Version	Supported Operating Systems	Supported Key Trustee Server Versions
3.16.x	<ul style="list-style-type: none"> • RHEL and CentOS: 7.7, 5.7, 5.10, 6.4 - 6.10, 7.1 - 7.5 • Oracle Linux: 7.7, 6.4 - 6.9, 7.1, 7.4, 7.5 • SLES: 11 SP2 - SP4, 12 SP1 - SP3 • Debian: 7.1, 7.8 • Ubuntu: 14.04, 14.04.3, 16.04 	5.4.x - 5.16.x
3.15.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.9, 7.1 - 7.5 • Oracle Linux: 6.4 - 6.9, 7.1, 7.4, 7.5 • SLES: 11 SP2 - SP4, 12 SP1 - SP3 • Debian: 7.1, 7.8 • Ubuntu: 12.04, 14.04, 14.04.3, 16.04 	5.4.x - 5.15.x
3.14.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.9, 7.1 - 7.4 • Oracle Linux: 6.4 - 6.9, 7.1, 7.4 • SLES: 11 SP2 - SP4; 12 SP1, SP2 • Debian: 7.1, 7.8 • Ubuntu: 12.04, 14.04, 14.04.3, 16.04 	5.4.x - 5.13.x
3.13.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.9, 7.1 - 7.4 • Oracle Linux: 6.4 - 6.9, 7.1 • SLES: 11 SP2 - SP4; 12 SP1, SP2 • Debian: 7.1, 7.8 • Ubuntu: 12.04, 14.04, 14.04.3, 16.04 	5.4.x - 5.12.x
3.12.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.9, 7.1 - 7.3 • Oracle Linux: 6.4 - 6.9, 7.1 • SLES: 11 SP2 - SP4; 12 SP1, SP2 • Debian: 7.1, 7.8 • Ubuntu: 12.04, 14.04, 14.04.3, 16.04 	5.4.x - 5.12.x
3.11.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.8, 7.1 - 7.3 • Oracle Linux: 6.4 - 6.7, 7.1 • SLES: 11 SP2 - SP4, 12 SP1 • Debian: 7.1, 7.8 • Ubuntu: 12.04, 14.04, 14.04.3 	5.4.x - 5.11.x

Cloudera Navigator Encrypt Version	Supported Operating Systems	Supported Key Trustee Server Versions
3.10.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.8, 7.1, 7.2 • Oracle Linux: 6.4 - 6.7, 7.1 • SLES: 11 SP2 - SP4 • Debian: 7.1, 7.8 • Ubuntu: 12.04, 14.04, 14.04.3 	5.4.x - 5.11.x
3.9.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.7, 7.1, 7.2 • Oracle Linux: 6.4 - 6.7, 7.1 • SLES: 11 SP2 - SP4 • Debian: 7.1, 7.8 • Ubuntu: 12.04, 14.04, 14.04.3 	5.4.x - 5.11.x
3.8.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.7, 7.1 • Oracle Linux: 6.4 - 6.7, 7.1 • SLES: 11 SP2, SP3 • Debian: 7.1 • Ubuntu: 12.04, 14.04 	5.4.x - 5.11.x
3.7.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 5.10, 6.4 - 6.6 • Oracle Linux: 6.4 - 6.6 • SLES: 11 SP2, SP3 • Debian: 7.1 • Ubuntu: 12.04, 14.04 	5.4.x - 5.11.x
3.6.x	<ul style="list-style-type: none"> • RHEL and CentOS: 5.7, 6.4, 6.5 • Oracle Linux: 6.4, 6.5 • SLES: 11 SP2, SP3 • Debian: 7.1 • Ubuntu: 12.04, 14.04 	3.8.x

Product Compatibility Matrix for Apache Sentry

Sentry enables role-based, fine-grained authorization for HiveServer2 and provides classic database-style authorization for Hive, Cloudera Impala and Cloudera Search. You can use either the Sentry service (introduced in Cloudera Manager 5.1.0 and CDH 5.1.0) or the policy file approach to secure your data.



Note: It's possible for a single cluster to use both the Sentry service (for Hive and Impala) and Sentry policy files (for Solr).

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version	Integrated into CDH
Apache Sentry	Sentry Service	Cloudera Manager 5.1.0	CDH 5.1.0	Impala 1.4.0 for CDH 5	Search for CDH 5.8.0	Yes
	Policy File	Cloudera Manager 5.0.0	CDH 5.0.0	Impala 1.3.0	Search 1.1.0	Yes

Support for Other Filesystems

Filesystems	Lowest Release Supported
Amazon S3 (s3a)	Apache Sentry for CDH 5.8
Amazon RDS	Apache Sentry for CDH 5.9

Product Compatibility Matrix for Apache Spark

Spark is a fast, general engine for large-scale data processing. For installation and configuration instructions, see [Setting Up Apache Spark Using the Command Line](#). To see new features introduced with each release, refer to the [CDH 5 Release Notes](#) on page 1812.

Since Spark 1.6 has been integrated into the CDH package, its compatibility with Cloudera Manager and CDH depends on the CDH 5.x.x release it is shipped with.



Note:

This page contains information related to Spark 1.6, which is included with CDH. For information about the separately available parcel for CDS 2 Powered by Apache Spark, see [the documentation for CDS 2](#).




Product	Lowest Supported Cloudera Manager Version	Lowest Supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version	Integrated into CDH
Apache Spark 1.6.x	Cloudera Manager 5.7.x	CDH 5.7.x	Not Supported	Not Supported	CDH 5.7.0
Apache Spark 1.5.x	Cloudera Manager 5.5.x and Cloudera Manager 5.6.x	CDH 5.5.x and CDH 5.6.x	Not Supported	Not Supported	CDH 5.5.0 and CDH 5.6.0
Apache Spark 1.3.x	Cloudera Manager 5.4.x	CDH 5.4.x	Not Supported	Not Supported	CDH 5.4.0
Apache Spark 1.2.x	Cloudera Manager 5.3.x	CDH 5.3.x	Not Supported	Not Supported	CDH 5.3.0
Apache Spark 1.1.x	Cloudera Manager 5.2.x	CDH 5.2.x	Not Supported	Not Supported	CDH 5.2.0
Apache Spark 1.0.x	Cloudera Manager 5.1.x	CDH 5.1.x	Not Supported	Not Supported	CDH 5.1.0

Product	Lowest Supported Cloudera Manager Version	Lowest Supported CDH Version	Lowest supported Impala Version	Lowest supported Search Version	Integrated into CDH
Apache Spark 0.9.x	Cloudera Manager 5.0.x	CDH 5.0.x	Not Supported	Not Supported	CDH 5.0.0

Product Compatibility for Dell EMC Isilon

See also [Hadoop Distributions and Products Supported by OneFS](#) on the Dell EMC Community website.

For compatibility with Cloudera Manager Replication and Snapshot features, see [Product Compatibility Matrix for Backup and Disaster Recovery](#) on page 77. For information on using Isilon with CDH, see [Using CDH with Isilon Storage](#).

CDH Version	Isilon OneFS Version
5.12 - 5.16 <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  Note: 5.12 and higher include Cloudera Navigator support to extract metadata and lineage from clusters using Isilon storage. </div>	<ul style="list-style-type: none"> • 8.1.2.x • 8.1.1.x • 8.1.0.x • 8.0.1.x • 8.0.0.x releases 8.0.0.4 and higher
5.8 - 5.11	<ul style="list-style-type: none"> • 8.1.2.x • 8.1.1.x • 8.1.0.x • 8.0.1.x • 8.0.0.x releases 8.0.0.4 and higher • 8.0.0.3 (with patch 184679) <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  Note: The patch fixes a bug that prevented high availability services from working correctly in a Kerberized environment. </div>
5.4.4 - 5.7	<ul style="list-style-type: none"> • 8.1.2.x • 8.1.1.x • 8.1.0.x • 8.0.1.x • 8.0.0.x releases 8.0.0.4 and higher • 8.0.0.3 (with patch 184679) <div style="border: 1px solid black; padding: 5px; margin-top: 10px;">  Note: The patch fixes a bug that prevented high availability services from working correctly in a Kerberized environment. </div> <ul style="list-style-type: none"> • 7.2.0.3 and higher (Cloudera recommends 7.2.1.1)

Product Compatibility Matrix for Backup and Disaster Recovery

This matrix contains compatibility information across features of Cloudera Manager Backup and Disaster Recovery and CDH and Cloudera Manager.

Product	Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Supported Services
Backup & Disaster Recovery	Replication	Cloudera Manager 4.5.0	CDH 4.0.1	HDFS, Hive, Impala
	Replication to and from Amazon S3*	Cloudera Manager 5.9.0	CDH 5.9.0	HDFS, Hive, Impala
	Replication to and from Microsoft ADLS	Cloudera Manager 5.15.0	CDH 5.15.0	HDFS, Hive, Impala
	Snapshots	Cloudera Manager 5.0.0	CDH 5.0.0	HDFS, Hive, Impala
	Snapshots from Isilon storage	Not Supported	Not Supported	HDFS, Hive, Impala

*BDR does not support S3 as a source or destination when S3 is configured to use SSE-KMS.

Supported Replication Scenarios

In Cloudera Manager 5, replication is supported between CDH 5 or CDH 4 clusters. The following tables describe support for HDFS and Hive/Impala replication. The versions listed in the table are the lowest version of Cloudera Manager and CDH required to perform the replication. For example, replication with TLS/SSL enabled on Hadoop services on the source and destination clusters requires Cloudera Manager and CDH version 5.0 or higher on the source and destination.

The tables below list the supported replication scenarios:



Note: BDR supports replication scenarios where TLS is enabled for non-Hadoop services (Hive/Impala) and TLS is disabled Hadoop services (such as HDFS, YARN, and MapReduce) if the Cloudera Manager and CDH versions are a supported replication scenario.

Service	Source			Destination		
	Cloudera Manager Version	CDH Version	Comment	Cloudera Manager Version	CDH Version	Comment
HDFS, Hive	4	4		5	4	
HDFS, Hive	4	4.4		5	5	
HDFS, Hive	5	5	TLS/SSL enabled on Hadoop services	5	5	TLS/SSL enabled on Hadoop services
HDFS, Hive	5	5	TLS/SSL enabled on Hadoop services	5	5	TLS/SSL <i>not</i> enabled on Hadoop services
HDFS, Hive	5	5.1	TLS/SSL enabled on Hadoop services and YARN	5	4 or 5	

CDH 5 and Cloudera Manager 5 Requirements and Supported Versions

Service	Source			Destination		
	Cloudera Manager Version	CDH Version	Comment	Cloudera Manager Version	CDH Version	Comment
HDFS, Hive	5	4		4.7.3 or higher	4	
HDFS, Hive	5	4		5	4	
HDFS, Hive	5	5		5	5	
HDFS, Hive	5	5		5	4.4 or higher	
HDFS, Hive	5.7	5.7, with Isilon storage	See Supported Replication Scenarios for Clusters using Isilon Storage on page 80.	5.7	5.7, with Isilon storage	See Supported Replication Scenarios for Clusters using Isilon Storage on page 80.
HDFS, Hive	5.7	5.7		5.7	5.7, with Isilon storage	
HDFS, Hive	5.7	5.7, with Isilon storage		5.7	5.7	
HDFS, Hive	5.8	5.8, with or without Isilon Storage	<ul style="list-style-type: none"> • Kerberos can be enabled • See Supported Replication Scenarios for Clusters using Isilon Storage on page 80. 	5.8	5.8, with or without Isilon Storage	<ul style="list-style-type: none"> • Kerberos can be enabled • See Supported Replication Scenarios for Clusters using Isilon Storage on page 80.

Unsupported Replication Scenarios



Note: If you are using Isilon storage for CDH, see [Supported Replication Scenarios for Clusters using Isilon Storage](#) on page 80.

The tables below list the unsupported replication scenarios:

Service	Source			Destination		
	Cloudera Manager Version	CDH Version	Comment	Cloudera Manager Version	CDH Version	Comment
Any	4 or 5	4 or 5	Kerberos enabled.	4 or 5	4 or 5	Kerberos not enabled
Any	4 or 5	4 or 5	Kerberos not enabled.	4 or 5	4 or 5	Kerberos enabled
HDFS, Hive	4 or 5	4	Where the replicated data includes a directory that contains a large number of files or subdirectories (several hundred thousand entries), causing out-of-memory errors.	4 or 5	5	

Service	Source			Destination		
	Cloudera Manager Version	CDH Version	Comment	Cloudera Manager Version	CDH Version	Comment
			To work around this issue, follow this procedure .			
Hive	4 or 5	4	Replicate HDFS Files is disabled.	4 or 5	4 or 5	Over-the-wire encryption is enabled.
Hive	4 or 5	4	Replication can fail if the NameNode fails over during replication.	4 or 5	5, with high availability enabled	Replication can fail if the NameNode fails over during replication.
Hive	4 or 5	4	The clusters use different Kerberos realms.	4 or 5	5	An older JDK is deployed. (Upgrade the CDH 4 cluster to use JDK 7 or JDK6u34 to work around this issue.)
Any	4 or 5	4	SSL enabled on Hadoop services.	4 or 5	4 or 5	
Hive	4 or 5	4.2 or higher	If the Hive schema contain views.	4 or 5	4	
HDFS	4 or 5	4, with high availability enabled	Replications fail if NameNode failover occurs during replication.	4 or 5	5, without high availability	Replications fail if NameNode failover occurs during replication.
HDFS	4 or 5	4 or 5	Over the wire encryption is enabled.	4 or 5	4	
HDFS	4 or 5	5	Clusters where there are URL-encoding characters such as % in file and directory names.	4 or 5	4	
Hive	4 or 5	4 or 5	Over the wire encryption is enabled and Replicate HDFS Files is enabled.	4 or 5	4	
Hive	4 or 5	4 or 5	From one cluster to the same cluster.	4 or 5	4 or 5	From one cluster to the same cluster.
HDFS, Hive	4 or 5	5	Where the replicated data includes a directory that contains a large number of files or subdirectories (several hundred thousand entries), causing out-of-memory errors. To work around this issue, follow this procedure .	4 or 5	4	
HDFS	4 or 5	5	The clusters use different Kerberos realms.	4 or 5	4	An older JDK is deployed. (Upgrade the CDH 4 cluster to use JDK 7 or JDK6u34

Service	Source			Destination		
	Cloudera Manager Version	CDH Version	Comment	Cloudera Manager Version	CDH Version	Comment
						to work around this issue.)
Hive	4 or 5	5	Replicate HDFS Files is enabled and the clusters use different Kerberos realms.	4 or 5	4	An older JDK is deployed. (Upgrade the CDH 4 cluster to use JDK 7 or JDK6u34 to work around this issue.)
Any	4 or 5	5	SSL enabled on Hadoop services and YARN.	4 or 5	4 or 5	
Any	4 or 5	5	SSL enabled on Hadoop services.	4 or 5	4	
HDFS	4 or 5	5, with high availability enabled	Replications fail if NameNode failover occurs during replication.	4 or 5	4, without high availability	Replications fail if NameNode failover occurs during replication.
HDFS, Hive	5	5		4	4	
Hive	5.2	5.2 or lower	Replication of Impala UDFs is skipped.	4 or 5	4 or 5	

Supported Replication Scenarios for Clusters using Isilon Storage

Note the following when scheduling replication jobs for clusters that use Isilon storage:

- As of CDH 5.8 and higher, Replication is supported for clusters using Kerberos and Isilon storage on the source or destination cluster, or both. See [Configuring Replication with Kerberos and Isilon](#). Replication between clusters using Isilon storage and Kerberos is not supported in CDH 5.7.
- Make sure that the `hdfs` user is a superuser in the Isilon system. If you specify alternate users with the **Run As** option when creating replication schedules, those users must also be superusers.
- Cloudera recommends that you use the Isilon `root` user for replication jobs. (Specify `root` in the **Run As** field when creating replication schedules.)
- Select the **Skip checksum checks** property when creating replication schedules.
- Clusters that use Isilon storage do not support [snapshots](#). Snapshots are used to ensure data consistency during replications in scenarios where the source files are being modified. Therefore, when replicating from an Isilon cluster, Cloudera recommends that you do not replicate Hive tables or HDFS files that could be modified before the replication completes.

See [Using CDH with Isilon Storage](#).

Supported Configurations with Virtualization and Cloud Platforms

This section lists supported cloud and virtualization platforms for deploying Cloudera software. The linked reference architectures are not replacements for statements of support above, but are guides to assist with deployment and sizing options.

Amazon Web Services

For information on deploying Cloudera software on a Amazon Web Services (AWS) cloud infrastructure, see the [Cloudera Enterprise Reference Architecture for AWS Deployments](#).

Google Cloud Platform

For information on deploying Cloudera software on a Google Cloud Platform infrastructure, see the [Cloudera Enterprise Reference Architecture for Google Cloud Platform Deployments](#).

Microsoft Azure

For information on deploying Cloudera software on a Microsoft Azure cloud infrastructure, see the [Cloudera Enterprise Reference Architecture for Azure Deployments](#).

Support Limitations for CDH and Cloudera Manager in Microsoft Azure

- Virtual machines must use Cloudera-published CentOS 6.x/7.x OS Images or Red Hat published 6.x/7.x Images. Red Hat images must include bootstrapped configuration as specified in the [Azure Bootstrap Scripts folder](#) in the Cloudera Director GitHub repository.
- Master Node Services *must use* Azure Premium Storage Disks for persistent storage.
- The VM local [temporary disk](#) *must not* be used for any persistent data for HDFS, Kudu, or other services.
- [Microsoft Azure Storage - Block Blob](#) (previously known as Windows Azure Storage Blob, accessed via wasb:// URI's) is only supported for backups using Hadoop's DistCP. Other services are not supported running directly against Azure Storage block blobs.

VMware

For information on deploying Cloudera software on a VMware-based infrastructure, see the [Reference architecture for deploying on VMware](#).

Recommendation when deploying on VMware in the current release:

- Use the part of Hadoop Virtual Extensions that has been implemented in [HADOOP-8468](#). This will prevent data loss when a physical node that hosts two or more DataNodes goes down .

Red Hat OpenStack Platform (OSP) 11

For information on deploying Cloudera software on RedHat OpenStack Platform (OSP) 11, see [Reference Architecture for Deploying CDH 5.x on Red Hat OSP 11](#).

For information on deploying Cloudera software on RedHat OpenStack Platform (OSP) 11 with Ceph storage, see [Reference Architecture for Deploying Cloudera Enterprise 5.x on Red Hat OpenStack Platform 11 with Red Hat Ceph Storage 2.x](#).

Hardware Requirements Guide

To assess the hardware and resource allocations for your cluster, you need to analyze the types of workloads you want to run on your cluster, and the CDH components you will be using to run these workloads. You should also consider the size of data to be stored and processed, the frequency of workloads, the number of concurrent jobs that need to run, and the speed required for your applications.

As you create the architecture of your cluster, you will need to allocate Cloudera Manager and CDH roles among the hosts in the cluster to maximize your use of resources. Cloudera provides some guidelines about how to assign roles to cluster hosts. See [Recommended Cluster Hosts and Role Distribution](#). When multiple roles are assigned to hosts, add together the total resource requirements (memory, CPUs, disk) for each role on a host to determine the required hardware.

For information about how workloads affect sizing decisions, see the following blog post: [How-to: Select the Right Hardware for Your New Hadoop Cluster](#).

For more information about sizing for a particular component, see the following minimum requirements:

Cloudera Manager

Cloudera Manager Server Storage Requirements

Component	Storage	Notes
Partition hosting /usr	1 GB	
Cloudera Manager Database	5 GB	If the Cloudera Manager Database shares a host with the Service Monitor and Host Monitor, more storage space is required to meet the requirements for those components.

Host Based Cloudera Manager Server Requirements

Number of Cluster Hosts	Database Host Configuration	Heap Size	Logical Processors	Cloudera Manager Server Storage Local Directory
Very small (≤ 10)	Shared	2 GB	4	5 GB minimum
Small (≤ 20)	Shared	4 GB	6	20 GB minimum
Medium (≤ 200)	Dedicated	8 GB	6	200 GB minimum
Large (≤ 500)	Dedicated	10 GB	8	500 GB minimum
Extra Large (> 500)	Dedicated	16 GB	16	1 TB minimum



Note: On smaller clusters, the Cloudera Manager Server and Database can share a host. On larger clusters, they must run on separate dedicated hosts.

Service Monitor Requirements

The requirements for the Service Monitor are based on the number of monitored entities. To see the number of monitored entities, perform the following steps:

1. Open the Cloudera Manager Admin Console and click **Clusters > Cloudera Management Service**.
2. Find the **Cloudera Management Service Monitored Entities** chart. If the chart does not exist, add it from the **Chart Library**.

For more information about Cloudera Manager entities, see [Cloudera Manager Entity Types](#).

Clusters with HDFS, YARN, or Impala

Use the recommendations in this table for clusters where the only services with worker roles are HDFS, YARN, or Impala.

Number of Monitored Entities	Number of Hosts	Required Java Heap Size	Recommended Non-Java Heap Size
0-2,000	0-100	1 GB	6 GB
2,000-4,000	100-200	1.5 GB	6 GB
4,000-8,000	200-400	1.5 GB	12 GB
8,000-16,000	400-800	2.5 GB	12 GB
16,000-20,000	800-1,000	3.5 GB	12 GB

Clusters with HBase, Solr, Kafka, or Kudu

Use these recommendations when services such as HBase, Solr, Kafka, or Kudu are deployed in the cluster. These services typically have larger quantities of monitored entities.

Number of Monitored Entities	Number of Hosts	Required Java Heap Size	Recommended Non-Java Heap Size
0-30,000	0-100	2 GB	12 GB
30,000-60,000	100-200	3 GB	12 GB
60,000-120,000	200-400	3.5 GB	12 GB
120,000-240,000	400-800	8 GB	20 GB

Host Monitor

The requirements for the Host Monitor are based on the number of monitored entities.

To see the number of monitored entities, perform the following steps:

1. Open the Cloudera Manager Admin Console and click **Clusters > Cloudera Management Service**.
2. Find the **Cloudera Management Service Monitored Entities** chart. If the chart does not exist, add it from the **Chart Library**.

For more information about Cloudera Manager entities, see [Cloudera Manager Entity Types](#).

Number of Hosts	Number of Monitored Entities	Heap Size	Non-Java Heap Size
0-200	<6k	1 GB	2 GB
200-800	6k-24k	2 GB	6 GB
800-1000	24k-30k	3 GB	6 GB

Ensure that you have at least 25 GB of disk space available for the Host Monitor, Service Monitor, Reports Manager, and Events Server databases.

For more information refer to [Host Monitor and Service Monitor Memory Configuration](#).

Reports Manager

The Reports Manager fetches the fsimage from the NameNode at regular intervals. It reads the fsimage and creates a Lucene index for it. To improve the indexing performance, Cloudera recommends provisioning a host as powerful as possible and dedicating an SSD disk to the Reports Manager.

Table 17: Reports Manager

Component	Java Heap	CPU	Disk
Reports Manager	3-4 times the size of the fsimage.	<ul style="list-style-type: none"> • Minimum: 8 cores • Recommended: 16 cores (32 cores, with hyperthreading enabled.) 	1 dedicated disk that is at least 20 times the size of the fsimage. Cloudera strongly recommends using SSD disks.

Agent Hosts

An unpacked parcel requires approximately three times the space of the packed parcel that is stored on the Cloudera Manager Server.

The /var/log partition on each host should have a minimum of 2GB of disk space allocated **per role** running on the host.

Event Server

The following table lists the minimum requirements for the Event Server:

CPU	RAM	Storage
1 core	256 MB	<ul style="list-style-type: none"> • 5 GB for the Event Database • 20 GB for the Event Server Index Directory. The location of this directory is set by the Event Server Index Directory Event Server configuration property.

Alert Publisher

The following table lists the minimum requirements for the Alert Publisher:


CPU	RAM	Storage
1 core	1 GB	Minimum of 1 disk for log files

Cloudera Navigator

The sizing of Navigator components varies heavily depending on the size of the cluster and the number of audit events generated. Refer to [Minimum Recommended Memory and Disk Space](#) for more information.

Table 18: Cloudera Navigator

Component	Java Heap / Memory	CPU	Disk
Navigator Audit Server	<p>Minimum: 2-3 GB of Java Heap</p> <p>Configure this value using the Java Heap Size of Auditing Server in Bytes configuration property.</p>	Minimum: 1 core	<p>The database used by the Navigator Audit Server must be able to accommodate hundreds of gigabytes (or tens of millions of rows per day). The database size may reach a terabyte.</p> <p>Ideally, the database should not be shared with other services because the audit insertion rate can overwhelm the database server making other services using same database less responsive.</p> <p>See Storage Space Planning for Cloudera Navigator.</p>
Navigator Metadata Server	<ul style="list-style-type: none"> • Minimum: 10 GB of Java Heap • Recommended: 20 GB Java Heap 	Minimum: 1 core	<ul style="list-style-type: none"> • Minimum: 50 GB • Recommended: 100-200 GB

Component	Java Heap / Memory	CPU	Disk
	<p>Add 20 GB for operating system buffer cache, however memory requirements can be much higher on a busy cluster and could require provisioning a dedicated host. Navigator logs include estimates based on the number of objects it is tracking.</p> <p>Configure this value using the Java Heap Size of Navigator Metadata Server in Bytes configuration property.</p>		<p>Note:  Data stored by the Metadata server grows indefinitely unless you run the purge function. If you have not set up the purge function to run at a scheduled interval, Cloudera recommends that you run the purge function to reclaim disk space and keep data growth (and the corresponding memory requirement) in check.</p> <p>See Administration (Navigator Console).</p>

Cloudera Data Science Workbench

Hardware Component	Requirement	Notes
CPU	16+ CPU (vCPU) cores	Allocate at least 1 CPU core per session. 1 CPU core is often adequate for light workloads.
Memory	32 GB RAM	<ul style="list-style-type: none"> As a general guideline, Cloudera recommends nodes with RAM between 60GB and 256GB Allocating less than 2 GB of RAM can lead to out-of-memory errors for many applications.
Disk	<ul style="list-style-type: none"> Root Volume: 100 GB 	SSDs are strongly recommended for application data storage.

Hardware Component	Requirement	Notes
	<ul style="list-style-type: none"> Application Block Device or Mount Point (Master Host Only): 1 TB Docker Image Block Device: 1 TB 	

For more information on scaling guidelines and storage requirements for cloud providers such as AWS and Azure, see [Requirements and Supported Platforms](#) in the Cloudera Data Science Workbench documentation.

CDH

Accumulo

Component	Java Heap	CPU	Disk
Master	<p>Minimum: 1 GiB</p> <ul style="list-style-type: none"> 100-1000 tablets: 4 GB 10,000 or more tablets with 200 or more Tablet Servers: 8 GB 10,000 or more tablets with 300 or more Tablet Servers: 12 GB <p>Set this value using the Tablet Server Max Heapsize Accumulo configuration property.</p> <p>See Flume Memory Consumption</p>	2 Cores. Add more for large clusters or bulk load.	1 disk for local logs
Tablet Server	<ul style="list-style-type: none"> Minimum: 512 MB Java Heap Recommended: 8 GB Medium scale production (200 or more tablets per Tablet Server): 16 GB <p>Set this value using the Tablet Server Max Heapsize Accumulo configuration property.</p>	Minimum 4 dedicated cores. You can add more cores for larger clusters, when using replication, or for bulk loads.	<ul style="list-style-type: none"> 4 or more spindles for each HDFS DataNode 1 disk for local logs (this disk can be shared with the operating system and/or other Hadoop logs)
Tracer	<p>1 - 2 GB, depending on cluster workloads</p> <p>Set this value using the Tracer Max Heapsize Accumulo configuration property.</p>	2 or more dedicated cores, depending on cluster size and workloads	1 disk for local logs, which can be shared with the operating system and/or other Hadoop logs

Component	Java Heap	CPU	Disk
GC Role	1-2 GB Set this value using the Garbage Collector Max Heapsize Accumulo configuration property.	2 or more dedicated cores, depending on cluster size	1 disk for local logs, which can be shared with the operating system and/or other Hadoop logs
Monitor Role	1-2 GB Set this value using the Monitor Max Heapsize Accumulo configuration property.	2 or more dedicated cores, depending on cluster size and workloads	1 disk for local logs, which can be shared with the operating system and/or other Hadoop logs

For additional information, see [Apache Accumulo on CDH Installation Guide](#).

Flume

Component	Java Heap	CPU	Disk
Flume	<ul style="list-style-type: none"> • Minimum: 1 GB • Maximum 4 GB • Java Heap size should be greater than the maximum channel capacity <p>Set this value using the Java Heap Size of Agent in Bytes Flume configuration property.</p> <p>See Flume Memory Consumption</p>	<p>Calculate the number of cores using the following formula:</p> $\frac{\text{Number of sources} + \text{Number of sinks}}{2}$	Multiple disks are recommended for file channels, either a JBOD setup or RAID10 (preferred due to increased reliability).

HDFS

Table 19: HDFS

Component	Memory	CPU	Disk
JournalNode	1 GB (default) Set this value using the Java Heap Size of JournalNode in Bytes HDFS configuration property.	1 core minimum	1 dedicated disk
NameNode	<ul style="list-style-type: none"> • Minimum: 1 GB (for proof-of-concept deployments) • Add an additional 1 GB for each additional 1,000,000 blocks <p>Snapshots and encryption can increase</p>	Minimum of 4 dedicated cores; more may be required for larger clusters	<ul style="list-style-type: none"> • Minimum of 2 dedicated disks for metadata • 1 dedicated disk for log files (This disk may be shared with the operating system.) • Maximum disks: 4

Component	Memory	CPU	Disk
	<p>the required heap memory.</p> <p>See Sizing NameNode Heap Memory</p> <p>Set this value using the Java Heap Size of NameNode in Bytes HDFS configuration property.</p>		
DataNode	<p>Minimum: 4 GB</p> <p>Increase the memory for higher replica counts or a higher number of blocks per DataNode. When increasing the memory, Cloudera recommends an additional 1 GB of memory for every 1 million replicas above 4 million on the DataNodes. For example, 5 million replicas require 5 GB of memory.</p> <p>Set this value using the Java Heap Size of DataNode in Bytes HDFS configuration property.</p>	<p>Minimum: 4 cores. Add more cores for highly active clusters.</p>	<p>Minimum: 4</p> <p>Maximum: 24</p> <p>The maximum acceptable size will vary depending upon how large average block size is. The DN's scalability limits are mostly a function of the number of replicas per DN, not the overall number of bytes stored. That said, having ultra-dense DNs will affect recovery times in the event of machine or rack failure. Cloudera does not support exceeding 100 TB per data node. You could use 12 x 8 TB spindles or 24 x 4TB spindles. Cloudera does not support drives larger than 8 TB.</p>




Warning: Running CDH on storage platforms other than direct-attached physical disks can provide suboptimal performance. Cloudera Enterprise and the majority of the Hadoop platform are optimized to provide high performance by distributing work across a cluster that can utilize data locality and fast local I/O. Refer to the [Cloudera Enterprise Storage Device Acceptance Criteria Guide](#) for more information about using non-local storage.

HBase

Component	Java Heap	CPU	Disk
Master	<ul style="list-style-type: none"> 100-10,000 regions: 4 GB 10,000 or more regions with 200 or more Region Servers: 8 GB 10,000 or more regions with 300 or more Region Servers: 12 GB 	<p>Minimum 4 dedicated cores. You can add more cores for larger clusters, when using replication, or for bulk loads.</p>	<p>1 disk for local logs, which can be shared with the operating system and/or other Hadoop logs</p>

Component	Java Heap	CPU	Disk
	Set this value using the Java Heap Size of HBase Master in Bytes HBase configuration property.		
Region Server	<ul style="list-style-type: none"> • Minimum: 8 GB • Medium-scale production: 16 GB • Heap larger than 16 GB requires special Garbage Collection tuning. See Configuring the HBase BlockCache <p>Set this value using the Java Heap Size of HBase RegionServer in Bytes HBase configuration property.</p>	Minimum: 4 dedicated cores	4 disks for each DataNode
Thrift Server	<p>1 GB - 4 GB</p> <p>Set this value using the Java Heap Size of HBase Thrift Server in Bytes HBase configuration property.</p>	Minimum 2 dedicated cores.	1 disk for local logs, which can be shared with the operating system and other Hadoop logs.

 **Note:** Consider adding more HBase Thrift Servers for production environments and deployments with a large number of Thrift client to scale horizontally.

Hive

Component	Java Heap	CPU	Disk	
HiveServer 2	Single Connection	4 GB	Minimum 4 dedicated cores	Minimum 1 disk This disk is required for the following: <ul style="list-style-type: none"> • HiveServer2 log files • <code>stdout</code> and <code>stderr</code> output files • Configuration files • Operation logs stored in the <code>operation_logs_dir</code> directory, which is configurable • Any temporary files that might be created by local map tasks
	2-10 connections	4-6 GB		
	11-20 connections	6-12 GB		
	21-40 connections	12-16 GB		
	41 to 80 connections	16-24 GB		
	Cloudera recommends splitting HiveServer2 into multiple instances and load balancing them once you start allocating more than 12 GB to HiveServer2. The objective is to adjust the size to reduce the impact of Java garbage collection on active processing by the service.			

Component	Java Heap	CPU	Disk
	<p>Set this value using the Java Heap Size of HiveServer2 in Bytes Hive configuration property.</p> <p>For more information, see Tuning Hive in CDH.</p>		under the /tmp directory
Hive Metastore	Single Connection	4 GB	Minimum 4 dedicated cores Minimum 1 disk This disk is required so that the Hive metastore can store the following artifacts: <ul style="list-style-type: none"> • Logs • Configuration files • Backend database that is used to store metadata if the database server is also hosted on the same node
	2-10 connections	4-10 GB	
	11-20 connections	10-12 GB	
	21-40 connections	12-16 GB	
	41 to 80 connections	16-24 GB	
	<p>Set this value using the Java Heap Size of Hive Metastore Server in Bytes Hive configuration property.</p> <p>For more information, see Tuning Hive in CDH.</p>		
Beeline CLI	Minimum: 2 GB	N/A	N/A

Hive on Spark Executor Nodes

Component	Memory	CPU	Disk
Hive-on-Spark	<ul style="list-style-type: none"> • Minimum: 16 GB • Recommended: 32 GB for larger data sizes <p>Individual executor heaps should be no larger than 16 GB so machines with more RAM can use multiple executors.</p>	<ul style="list-style-type: none"> • Minimum: 4 cores • Recommended: 8 cores for larger data sizes 	Disk space requirements are driven by scratch space requirements for Spark spill.
	<p>For more information on how to reserve YARN cores and memory that will be used by Spark executors, refer to Tuning Apache Hive on Spark in CDH.</p>		

HSM KMS

Component	Memory	CPU	Disk
Navigator HSM KMS	16 GB RAM	Minimum: 2 GHz 64-bit quad core	40 GB, using moderate to high-performance drives.

Hue

Component	Memory	CPU	Disk
Hue Server	<ul style="list-style-type: none"> • Minimum: 4 GB • Maximum 10 GB • If the cluster uses the Hue load balancer, add additional memory 	<p>Minimum: 1 Core to run Django</p> <p>When Hue is configured for high availability, add additional cores</p>	<p>Minimum: 10 GB for the database, which grows proportionally according to the cluster size and workloads.</p> <p>When Hue is configured for high availability, add temp space is required</p>



Note: Hue is limited by cgroup settings. Cloudera Manager 5.9 and higher all memory soft/hard limits are set to -1.

For more information about Hue high availability, see [How to Add a Hue Load Balancer](#).

Impala

Sizing requirements for Impala can vary significantly depending on the size and types of workloads using Impala.

Component	Native Memory	JVM Heap	CPU	Disk
Impala Daemon	<p>Set this value using the Impala Daemon Memory Limit configuration property.</p> <ul style="list-style-type: none"> • Minimum: 32 GB • Recommended: 128 GB 	<p>Use the <code>JAVA_TOOL_OPTIONS</code> environment variable to set the maximum heap size for the Coordinator Impala Daemons.</p> <ul style="list-style-type: none"> • Minimum: 4 GB • Recommended: 8 GB 	<ul style="list-style-type: none"> • Minimum: 4 • Recommended: 16 or more <p>CPU instruction set: AVX2</p>	<ul style="list-style-type: none"> • Minimum: 1 disk • Recommended: 8 or more
Catalog Server	<p>Set the maximum heap size for the Catalog Server:</p> <ul style="list-style-type: none"> • Minimum: 4 GB • Recommended: 8 GB <p>In Cloudera Manager, use the Java Heap Size of Catalog Server in Bytes configuration property (Cloudera Manager 5.7 and higher), or Impala Catalog Server Environment Advanced Configuration Snippet (Safety Valve) (Cloudera Manager 5.6 and lower).</p> <p>If not using Cloudera Manager, use the <code>JAVA_TOOL_OPTIONS</code> environment variable. For example, to set it to 8 GB: <code>JAVA_TOOL_OPTIONS= "-Xmx8g"</code></p>		<ul style="list-style-type: none"> • Minimum: 4 • Recommended: 16 or more <p>CPU instruction set: AVX2</p>	<ul style="list-style-type: none"> • Minimum and Recommended: 1 disk

For the networking topology for multi-rack cluster, [Leaf-Spine](#) is recommended for the optimal performance.

Kafka

Kafka requires a fairly small amount of resources, especially with some configuration tuning. By default, Kafka, can run on as little as 1 core and 1GB memory with storage scaled based on requirements for data retention.

CPU is rarely a bottleneck because Kafka is I/O heavy, but a moderately-sized CPU with enough threads is still important to handle concurrent connections and background tasks.

Kafka brokers tend to have a similar hardware profile to HDFS data nodes. How you build them depends on what is important for your Kafka use cases. Use the following guidelines:


To affect performance of these features:	Adjust these parameters:
Message Retention	Disk size
Client Throughput (Producer & Consumer)	Network capacity
Producer throughput	Disk I/O
Consumer throughput	Memory

A common choice for a Kafka node is as follows:


Component	Memory/Java Heap	CPU	Disk
Broker	<ul style="list-style-type: none"> RAM: 64 GB Recommended Java heap: 4 GB Set this value using the Java Heap Size of Broker Kafka configuration property. See	12- 24 cores	<ul style="list-style-type: none"> 1 HDD For operating system 1 HDD for Zookeeper dataLogDir 10- HDDs, using Raid 10, for Kafka data
MirrorMaker	1 GB heap Set this value using the Java Heap Size of MirrorMaker Kafka configuration property.	1 core per 3-4 streams	No disk space needed on MirrorMaker instance. Destination brokers should have sufficient disk space to store the topics being copied over.

Networking requirements: Gigabit Ethernet or 10 Gigabit Ethernet

Key Trustee Server

Component	Memory	CPU	Disk
Key Trustee Server <div style="border: 1px solid black; padding: 5px; width: fit-content;">  Note: KTS requires a additional dedicated resources. For more information, refer to Data at Rest Encryption Requirements. </div>	8 GB	1 GHz 64-bit quad core	20 GB, using moderate to high-performance drives

Key Trustee KMS

Component	Memory	CPU	Disk
Key Trustee KMS	16 GB	2 GHz 64-bit quad core	40 GB, using moderate to high-performance drives
<div style="border: 1px solid black; padding: 5px; margin: 5px 0;"> <p>Note:  Cloudera recommends using machines with capabilities equivalent to your NameNode hosts, with Intel CPUs that support AES-NI for optimum performance.</p> </div>			

Kudu

Component	Memory	CPU	Disk
Tablet Server	<ul style="list-style-type: none"> • Minimum: 4 GB • Recommended: 10 GB <p>Additional hardware may be required, depending on the workloads running in the cluster. If you are using Impala, see the Impala sizing guidelines.</p>		1 disk for write-ahead log (WAL). Using an SSD drive may improve performance.
Master	<ul style="list-style-type: none"> • Minimum: 256 MB • Recommended: 1 GB 		1 disk

For more information, see [Kudu Server Management](#).

Oozie

Component	Java Heap	CPU	Disk
Oozie	<ul style="list-style-type: none"> • Minimum: 1 GB (this is the default set by Cloudera Manager). This is sufficient for less than 10 simultaneous workflows, without forking. • If you notice excessive garbage collection, or out-of-memory errors, increase the heap size to 4 GB for 	No resources required	No resources required

Component	Java Heap	CPU	Disk
	medium-size production clusters or to 8 GB for large-size production clusters. <ul style="list-style-type: none"> Set this value using the Java Heap Size of Oozie Server in Bytes Oozie configuration property. 		

Additional tuning:

For workloads with many coordinators that run with complex workflows (a **max concurrency reached!** warning appears in the log and the Oozie `admin -queuedump` command shows a large queue):

- Increase the value of the `oozie.service.CallableQueueService.callable.concurrency` property to 50.
- Increase the value of the `oozie.service.CallableQueueService.threads` property to 200.

Do not use a Derby database as a backend database for Oozie.

Search

Component	Java Heap	CPU	Disk
Solr	<ul style="list-style-type: none"> Small workloads, or evaluations: 16 GB Smaller production environments: 32 GB Larger production environments: 96 GB is sufficient for most clusters. Set this value using the Java Heap Size of Solr Server in Bytes Solr configuration property. See	<ul style="list-style-type: none"> Minimum: 4 Recommended: 16 for production workloads 	No requirement. Solr uses HDFS for storage.

Note the following considerations for determining the optimal amount of heap memory:

- Size of searchable material:** The more searchable material you have, the more memory you need. All things being equal, 10 TB of searchable data requires more memory than 1 TB of searchable data.
- Content indexed in the searchable material:** Indexing all fields in a collection of logs, email messages, or Wikipedia entries requires more memory than indexing only the Date Created field.
- The level of performance required:** If the system must be stable and respond quickly, more memory may help. If slow responses are acceptable, you may be able to use less memory.

For more information refer to [Deployment Planning for Cloudera Search](#).

Sentry

Component	Java Heap	CPU	Disk
Sentry Server	<ul style="list-style-type: none"> Minimum: 576 MB 	Minimum: 4	

Component	Java Heap	CPU	Disk
	<ul style="list-style-type: none"> Recommended: 2.5 GB per million objects in the Hive database (servers, databases, tables, partitions, columns, URIs, and views) <p>Set this value with the Java Heap Size of Sentry Server in Bytes Sentry configuration property.</p>		

For more information about Sentry requirements, see [Before You Install Sentry](#).

Spark

Component	Java Heap	CPU	Disk
Spark History Server	<p>Minimum: 512 MB</p> <p>Set this value using the Java Heap Size of History Server in Bytes Spark configuration property.</p>	<p>1</p> <div style="border: 1px solid orange; padding: 5px; margin-top: 10px;"> <p>Important: Cloudera recommends that you adjust the number of CPUs and memory for the Spark History Server based on your specific cluster usage patterns.</p> </div>	<p>Minimum 1 disk for log files.</p>

YARN

Component	Java Heap	CPU	Other Recommendations
Job History Server	<ul style="list-style-type: none"> Minimum: 1 GB Increase memory by 1.6 GB for each 100,000 tasks kept in memory. For example: 5 jobs @ 100,000 mappers + 20,000 reducers = 600,000 total tasks requiring 9.6 GB of heap. <p>See the Other Recommendations column for additional tuning suggestions.</p> <p>Set this value using the Java Heap Size of JobHistory</p>	<p>Minimum: 1 core</p>	<ul style="list-style-type: none"> Set the mapreduce.jobhistory.hist.format property to binary (history files will load about 2x-3x faster with this setting). Available in CDH 5.5.0 or higher only. Set the mapreduce.jobhistory.loadtablesize property to a total loaded task count. Using the example in the Java Heap column to the left, of 650,000 total tasks, you can set it to 700,000 to allow for some safety margin.

Component	Java Heap	CPU	Other Recommendations
	<p>Server in Bytes YARN configuration property.</p>		<p>This should also prevent the JobHistoryServer from hanging during garbage collection, since the job count limit does not have a task limit.</p> <ul style="list-style-type: none"> It is strongly recommended to enable G1GC if the heap is large (larger than 10 GB). It is strongly recommended to set <code>mapreduce.jobhistory.maxjobspernode</code> so that jobs that have too many tasks will not be loaded and cached in memory. Such jobs can be retrieved by using the command line: <code>mapred job -history <jobId></code>
NodeManager	<p>Minimum: 1 GB</p> <p>Configure additional heap memory for the following conditions:</p> <ul style="list-style-type: none"> Large number of containers Large shuffle sizes in Spark or MapReduce <p>Set this value using the Java Heap Size of NodeManager in Bytes YARN configuration property.</p>	<ul style="list-style-type: none"> Minimum: 8-16 cores Recommended: 32-64 cores 	<p>Disks:</p> <ul style="list-style-type: none"> Minimum: 8 disks Recommended: 12 or more disks <p>Networking:</p> <ul style="list-style-type: none"> Minimum: Dual 1Gbps or faster Recommended: Single/Dual 10 Gbps or faster
ResourceManager	<p>Minimum: 1 GB</p> <p>Configure additional heap memory for the following conditions:</p> <ul style="list-style-type: none"> More jobs Larger cluster size Number of retained finished applications (configured with the <code>yarn.resourcemanager.completedapplications</code> property). Scheduler configuration <p>Set this value using the Java Heap Size of</p>	Minimum: 1 core	

Component	Java Heap	CPU	Other Recommendations
	ResourceManager in Bytes YARN configuration property.		
Other Settings	<ul style="list-style-type: none"> Set the ApplicationMaster Memory YARN configuration property to 512 MB Set the Container Memory Minimum YARN configuration property to 1 GB. 	N/A	N/A

For more information, see [Tuning YARN](#).

ZooKeeper

Component	Java Heap	CPU	Disk
ZooKeeper Server	<ul style="list-style-type: none"> Minimum: 1 GB Increase heap size when watching 10,000 - 100,000 ephemeral znodes and are using 1,000 or more clients. <p>Set this value using the Java Heap Size of ZooKeeper Server in Bytes ZooKeeper configuration property.</p>	Minimum: 4 cores	ZooKeeper was not designed to be a low-latency service and does not benefit from the use of SSD drives. The ZooKeeper access patterns – append-only writes and sequential reads – were designed with spinning disks in mind. Therefore Cloudera recommends using HDD drives.

Additional information:

- [Managing ZooKeeper](#)
- [ZooKeeper Administration Guide: Things to Avoid](#)

Version and Download Information

Version and download information for Cloudera Manager, CDH, Impala, and Search can be found in the HTML documentation on the website at [Cloudera Documentation](#). Select the release version number and go to the HTML version of the Release Guide.

Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication

Access to Cloudera Manager binaries for production purposes requires authentication. To access the binaries at the locations below, you must first have an active subscription agreement and obtain a license key file along with the required authentication credentials (username and password).

If your deployment was not installed using authenticated downloads and if you need to add hosts to an existing cluster, or add new clusters, you must follow the steps in this topic to configure your deployment to access the downloads.

Note that as of February 1, 2021, the download URLs on `archive.cloudera.com` have changed. They are the same as the previous URLs, except that `/p` is added to the URL, after `archive.cloudera.com`. Depending on where you use the URL, you may also need to prepend your credentials to the URL. The instructions below contain that information. For example:

```
https://username:password@archive.cloudera.com/p/...
```

For a complete list of download URLs, see:

- [Cloudera Manager 5.x Version and Download Information](#)
- [CDH 5.x Download Information](#)

Add hosts to an existing cluster

See

- [Add hosts to an existing cluster managed by Cloudera Manager using parcels](#) on page 102, below.
- [Add hosts to an existing cluster managed by Cloudera Manager using packages](#) on page 103

Add a new cluster

If you are adding a CDH cluster that is the same version of CDH as another cluster managed by the same instance of Cloudera Manager, no further action is required. To add a new cluster with a different supported version of CDH, see

- [Adding a Cluster to an existing Cloudera Manager deployment using Parcels](#) on page 103, below.
- [Adding a Cluster to an existing Cloudera Manager deployment using Packages](#) on page 104

New Installations

You must use the new URLs and include your username and password to access downloads for a new installation. See [Cloudera Installation Guide](#).

Upgrades

You must use the new URLs and include your username and password to access downloads for the upgrade. See the [Cloudera Enterprise Upgrade Guide](#).

If your deployment was installed or upgraded using authenticated downloads, no further action is needed.

Obtain Credentials

The license key file and authentication credentials are provided in an email sent to customer contacts from Cloudera when a new license is issued. If you have an existing license with a CDP Private Cloud Base Edition entitlement, you

Version and Download Information

might not have received an email. In this instance you can identify the authentication credentials from the license key file. If you do not have access to the license key, contact your account representative to receive a copy.

The license key file and authentication credentials are provided in an email sent to customer accounts from Cloudera when a new license is issued. If you have an existing license with a Cloudera Enterprise entitlement, you might not have received an email. In this instance you can identify the authentication credentials from the license key file. If you do not have access to the license key, contact your account representative to receive a copy.

To identify your authentication credentials using your license key file, complete the following steps:

1. From cloudera.com, log into the cloudera.com account associated with the Cloudera Enterprise license and subscription agreement.
2. On the [Cloudera CDH Download](#) page, select the version that you want to install and click **Get It Now!**
3. Scroll down to the Credentials Generator at the bottom of the page.
4. In the Generate Credentials text box, copy and paste the text of the “PGP Signed Message” within your license key file and click **Get Credentials**. The credentials generator returns your username and password.



Note: Take note of the authentication credentials. You will need them during installation to complete tasks such as configuring a remote or local package repository.

When you obtain your authentication credentials, use them to form the URLs below, where you can access the Cloudera repositories in the Cloudera Archive.

Prepare your deployment

There are two options for access to the downloads:

- **Option A: Upgrade Cloudera Manager and configure the package manager**
 1. **Upgrade Cloudera Manager** to one of the versions listed below. (If you are using Cloudera Manager 6.3.3 or higher, you do not need to upgrade.) Cloudera Manager is compatible with any version of CDH that is equal to lower than the Cloudera Manager version. The versions of Cloudera Manager shown below have been updated to use authenticated downloads:

Table 20: Cloudera Manager Upgrades for authenticated downloads

Current Cloudera Manager version	Upgrade Cloudera Manager to this version
5.0.0 - 5.16.2	5.16.2.4505



Note: You do not need to upgrade CDH.

To upgrade Cloudera Manager:

1. Go to the interactive [Cloudera Enterprise Upgrade guide page for upgrading Cloudera Manager](#).
 2. In the form at the top of the page, select the current version of your Cloudera Manager deployment.
 3. Select the appropriate new version of Cloudera Manager from the table above.
 - a. Select the operating system, database, and Navigator status.
 4. Follow the remaining steps on the page to upgrade the Cloudera Manager server and agents.
2. **Configure your operating system package manager** (yum, zypper, or apt-get) to access the public Cloudera repositories at archive.cloudera.com. See the following links;
- [Configure a Repository for Cloudera Manager 5.16](#)

- [Configure a Repository for Cloudera Manager 5.15](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.14](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.13](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.12](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.11](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.10](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.9](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.8](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.7](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.6](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.5](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.4](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.3](#)
 - [Establish Your Cloudera Manager Repository Strategy using Cloudera Manager 5.2](#)
- **Option B: Use a local package and/or parcel repository**
- **Use a local package and /or parcel repository.** If your cluster does not have access to the internet, this is the only option.
 1. Download the contents of the repository for your version of the software from `archive.cloudera.com` using your username and password.
 2. Do one of the following:
 - Create a [local package repository](#) (for Cloudera Manager and CDH clusters installed using packages). See the following links:
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.16](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.15](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.14](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.13](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.12](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.11](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.10](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.9](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.8](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.7](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.6](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.5](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.4](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.3](#)
 - [Configuring a Local package repository for Cloudera Manager/CDH 5.2](#)
 - Create *local parcel repository* (for CDH clusters installed using Parcels). See the following links:
 - [Configuring a Local parcel repository for Cloudera Manager/CDH 5.16](#)
 - [Configuring a Local parcel repository for Cloudera Manager/CDH 5.15](#)
 - [Configuring a Local parcel repository for Cloudera Manager/CDH 5.14](#)
 - [Configuring a Local parcel repository for Cloudera Manager/CDH 5.13](#)
 - [Configuring a Local parcel repository for Cloudera Manager/CDH 5.12](#)
 - [Configuring a Local parcel repository for Cloudera Manager/CDH 5.11](#)
 - [Configuring a Local parcel repository for Cloudera Manager/CDH 5.10](#)
 - [Configuring a Local parcel repository for Cloudera Manager/CDH 5.9](#)

- [Configuring a Local parcel repository for Cloudera Manager/CDH 5.8](#)
- [Configuring a Local parcel repository for Cloudera Manager/CDH 5.7](#)
- [Configuring a Local parcel repository for Cloudera Manager/CDH 5.6](#)
- [Configuring a Local parcel repository for Cloudera Manager/CDH 5.5](#)
- [Configuring a Local parcel repository for Cloudera Manager/CDH 5.4](#)
- [Configuring a Local parcel repository for Cloudera Manager/CDH 5.3](#)
- [Configuring a Local parcel repository for Cloudera Manager/CDH 5.2](#)

Add hosts to an existing cluster managed by Cloudera Manager using parcels

1. Follow the steps to *add hosts using the Add Host Wizard*. See the following links:

- [Adding a Host to the cluster using Cloudera Manager 5.16](#)
- [Adding a Host to the cluster using Cloudera Manager 5.15](#)
- [Adding a Host to the cluster using Cloudera Manager 5.14](#)
- [Adding a Host to the cluster using Cloudera Manager 5.13](#)
- [Adding a Host to the cluster using Cloudera Manager 5.12](#)
- [Adding a Host to the cluster using Cloudera Manager 5.11](#)
- [Adding a Host to the cluster using Cloudera Manager 5.10](#)
- [Adding a Host to the cluster using Cloudera Manager 5.9](#)
- [Adding a Host to the cluster using Cloudera Manager 5.8](#)
- [Adding a Host to the cluster using Cloudera Manager 5.7](#)
- [Adding a Host to the cluster using Cloudera Manager 5.5](#)
- [Adding a Host to the cluster using Cloudera Manager 5.4](#)
- [Adding a Host to the cluster using Cloudera Manager 5.3](#)
- [Adding a Host to the cluster using Cloudera Manager 5.2](#)

2. When you get to the **Select Repository** page, select **Use Parcels**.

3. Select the version of CDH that you want to install. If you want to use a version of CDH that is not displayed:

a. Click **More Options**.

b. In the **Remote Parcel Repository URLs** section, click the + icon do one of the following:

- If you have downloaded the parcel repository and have set up a local repository server, enter the URL for the local repository.
- If the cluster has access to the internet, enter the URL to the public repository on `archive.cloudera.com` See [CDH 5.x Download Information](#) to obtain the URL. For example:

```
https://username:password@archive.cloudera.com/p/cdh/parcels/5.16.2
```

The available versions of CDH display.

c. Select the version of CDH you want to install.

4. In the **Select the specific release of the Cloudera Manager Agent you want to install on your hosts** section, click **Custom Repository**.

5. Enter the URL for the Cloudera Manager package repository. For example:

```
https://<username>:<password>@archive.cloudera.com/p/cm5/redhat/6/x86_64/cm/5.16.2.4505
```

6. Enter the URL for the GPG signing key in the **Custom GPG Key URL** field.

For example:

```
https://<username>:<password>@archive.cloudera.com/p/cm5/redhat/6/x86_64/cm/RPM-GPG-KEY-cloudera
```

7. Continue with Add Hosts wizard to install the Cloudera Manager Agent, and, optionally the JDK.

Add hosts to an existing cluster managed by Cloudera Manager using packages

1. Follow the steps to *add hosts using the Add Host Wizard*. See the following links:

- [Adding a Host to the cluster using Cloudera Manager 5.16](#)
- [Adding a Host to the cluster using Cloudera Manager 5.15](#)
- [Adding a Host to the cluster using Cloudera Manager 5.14](#)
- [Adding a Host to the cluster using Cloudera Manager 5.13](#)
- [Adding a Host to the cluster using Cloudera Manager 5.12](#)
- [Adding a Host to the cluster using Cloudera Manager 5.11](#)
- [Adding a Host to the cluster using Cloudera Manager 5.10](#)
- [Adding a Host to the cluster using Cloudera Manager 5.9](#)
- [Adding a Host to the cluster using Cloudera Manager 5.8](#)
- [Adding a Host to the cluster using Cloudera Manager 5.7](#)
- [Adding a Host to the cluster using Cloudera Manager 5.5](#)
- [Adding a Host to the cluster using Cloudera Manager 5.4](#)
- [Adding a Host to the cluster using Cloudera Manager 5.3](#)
- [Adding a Host to the cluster using Cloudera Manager 5.2](#)

2. When you get to the **Select Repository** page, select **Use Packages**.

3. Select **Custom Repository** and enter one of the following:

- If you have downloaded the repository and set up a local repository server, enter URL for the local package repository.
- If the cluster has access to the internet, enter the URL to the public Cloudera package repository on `archive.cloudera.com`. See [Cloudera Manager Version and Download Information](#) to obtain the URL.

For example:

```
https://<username>:<password>@archive.cloudera.com/p/cm5/redhat/7/x86_64/cm/5.16.2/
```

4. In the **Select the specific release of the Cloudera Manager Agent you want to install on your hosts.** section, click **Custom Repository**.

5. Enter the URL for the Cloudera Manager package repository. For example:

```
https://<username>:<password>@archive.cloudera.com/p/cm5/redhat/6/x86_64/cm/5.16.2.4505
```

6. Enter the URL for the GPG signing key in the **Custom GPG Key URL** field.

For example:

```
https://<username>:<password>@archive.cloudera.com/p/cm5/redhat/6/x86_64/cm/RPM-GPG-KEY-cloudera
```

7. Continue with Add Hosts wizard to install the Cloudera Manager Agent, and, optionally the JDK.

Adding a Cluster to an existing Cloudera Manager deployment using Parcels

1. Follow the steps to *add a cluster*. See the following links:

- [Adding and Deleting Clusters using Cloudera Manager 5.16](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.15](#)

Version and Download Information

- [Adding and Deleting Clusters using Cloudera Manager 5.14](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.13](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.12](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.11](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.10](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.9](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.8](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.7](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.5](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.4](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.3](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.2](#)

2. Select **Use Parcels**.

The available versions of CDH display.

3. Select the version of CDH that you want to install. If you want to use a version of CDH that is not displayed:

a. Click **More Options**.

b. In the **Remote Parcel Repository URLs** section, click the + icon do one of the following:

- If you have downloaded the parcel repository and have set up a local repository server, enter the URL for the local repository.
- If the cluster has access to the internet, enter the URL to the public repository on `archive.cloudera.com` See [CDH 5.x Download Information](#) to obtain the URL. For example:

```
https://username:password@archive.cloudera.com/p/cdh/parcels/5.16.2
```

The available versions of CDH display.

c. Select the version of CDH you want to install.

4. In the **Select the specific release of the Cloudera Manager Agent you want to install on your hosts.** section, click **Custom Repository**.

5. Enter the URL for the Cloudera Manager package repository. For example:

```
https://<username>:<password>@archive.cloudera.com/p/cm5/redhat/6/x86_64/cm/5.16.2.4505
```

6. Enter the URL for the GPG signing key in the **Custom GPG Key URL** field.

For example:

```
https://<username>:<password>@archive.cloudera.com/p/cm5/redhat/6/x86_64/cm/RPM-GPG-KEY-cloudera
```

7. Continue with add cluster wizard to install the Cloudera Manager Agent, and, optionally the JDK.

Adding a Cluster to an existing Cloudera Manager deployment using Packages

1. Follow the steps to *add a cluster*. See the following links:

- [Adding and Deleting Clusters using Cloudera Manager 5.16](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.15](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.14](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.13](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.12](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.11](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.10](#)

- [Adding and Deleting Clusters using Cloudera Manager 5.9](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.8](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.7](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.5](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.4](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.3](#)
- [Adding and Deleting Clusters using Cloudera Manager 5.2](#)

2. Select **Use Packages**.

3. Select **Custom Repository**

4. Enter the custom repository URL.

- If you have downloaded the package repository and have set up a local repository server, enter URL for the local repository.
- If the cluster has access to the internet, enter the URL to the public Cloudera package repository on `archive.cloudera.com`. See [CDH Download Information](#) to obtain the URL.

For example:

```
https://<username>:<password>@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.16.2/
```

5. In the **Select the specific release of the Cloudera Manager Agent you want to install on your hosts.** section, click **Custom Repository**.

6. Enter the URL for the Cloudera Manager package repository. For example:

```
https://<username>:<password>@archive.cloudera.com/p/cm5/redhat/6/x86_64/cm/5.16.2.4505
```

7. Enter the URL for the GPG signing key in the **Custom GPG Key URL** field.

For example:

```
https://<username>:<password>@archive.cloudera.com/p/cm5/redhat/6/x86_64/cm/RPM-GPG-KEY-cloudera
```

8. Continue with add cluster wizard to install the Cloudera Manager Agent, and, optionally the JDK.

Cloudera Manager Version and Download Information

Cloudera Manager is available in the following releases:

- Cloudera Manager 5.16.2 is the current release of Cloudera Manager 5.16.
- Cloudera Manager 5.15.2, 5.14.4, 5.13.3, 5.12.2, 5.11.2, 5.10.2, 5.9.3, 5.8.5, 5.7.6, 5.6.1, 5.5.6, 5.4.10, 5.3.10, 5.2.7, 5.1.6, and 5.0.7 are previous stable releases of Cloudera Manager 5.14, 5.13, 5.12, 5.11, 5.10, 5.9, 5.8, 5.7, 5.6, 5.5, 5.4, 5.3, 5.2, 5.1, and 5.0 respectively.



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

**Warning:**

- Cloudera Manager 4 and CDH 4 have reached End of Maintenance (EOM) on August 9, 2015. Cloudera does not support or provide updates for Cloudera Manager 4 and CDH 4 releases.
- Cloudera Manager 3 and CDH 3 have reached End of Maintenance (EOM) on June 20, 2013. Cloudera does not support or provide updates for Cloudera Manager 3 and CDH 3 releases.

The 64-bit packages listed here support both Cloudera Express with its extensive set of monitoring and management features, and Cloudera Enterprise with additional functionality. A 60-day trial can be enabled to provide access to the full set of Cloudera Enterprise Cloudera Enterprise features. Cloudera Enterprise can be enabled permanently with the appropriate license. To obtain a Cloudera Enterprise license, fill in this [form](#) or call 866-843-7207.



Important: The repo file links in the Repositories tables in this topic always install the *latest* version of Cloudera Manager 5. If you need to install a specific, lower version of Cloudera Manager 5, you must download the contents of the repository from the "repo-as-tarball" directory and create an [internal package repository](#) to use for your installation. Open one of the following links, and then navigate to the correct operating system and Cloudera version:

- – **Cloudera Manager 5:**
<https://username:password@archive.cloudera.com/p/cm5/repo-as-tarball/>
- **CDH 5:** <https://username:password@archive.cloudera.com/p/cdh5/repo-as-tarball/>

Component Versions

For a complete list of the individual component versions, see the `manifest.json` file in the `parcel` directory for your chosen release. The component version numbers are the same in parcels as the package distributions for each release.

1. Go to <https://username:password@archive.cloudera.com/p/cdh5/parcels/>.
2. Open the directory for your release (for example, `/cdh5/parcels/5.9.0`).
3. Open the `manifest.json` file.

For each component, the `pkg_version` is a concatenation of `<upstream_version>+<cdh_version>+<number_of_patches_on_upstream_release>`.

For example, in `/cdh5/parcels/5.9.0/manifest.json`, this entry for Hadoop shows an upstream version is 2.6.0, the CDH `version is 5.9.0, and there have been 1,799 patches to the upstream version in the 2.6.0 release.

```
{
  "pkg_version": "2.6.0+cdh5.9.0+1799",
  "pkg_release": "1.cdh5.9.0.p0.30",
  "name": "hadoop",
  "version": "2.6.0-cdh5.9.0"
}
```

Cloudera Manager 5.16.2.4505

Release Date: February 1, 2021 **Status:** Production



Note: This version of Cloudera Manager has been modified to manage clusters using the new download locations that require authentication.

Documentation:

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Table 21: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/5.16.2.4505/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/x86_64/cm/5.16.2.4505.tar.gz
yum RHEL/CentOS/Oracle 6	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/5.16.2.4505/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/x86_64/cm/5.16.2.4505.tar.gz
yum RHEL/CentOS/Oracle 5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/5.16.2.4505/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/x86_64/cm/5.16.2.4505.tar.gz
apt Debian Jessie 8	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/dists/jessie-cm5.16.2.4505/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/jessie/amd64/cm/5.16.2.4505.tar.gz
apt Debian Wheezy 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/dists/wheezy-cm5.16.2.4505/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/wheezy/amd64/cm/5.16.2.4505.tar.gz
zypper/YaST SLES 12	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/5.16.2.4505/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/cloudera-cm.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/sles/12/x86_64/cm/5.16.2.4505.tar.gz
zypper/YaST SLES 11	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/5.16.2.4505/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/sles/11/x86_64/cm/5.16.2.4505.tar.gz
Ubuntu Xenial 16.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/xenial/amd64/cm/dists/xenial-cm5.16.2.4505/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/xenial/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/ubuntu/xenial/amd64/cm/5.16.2.4505.tar.gz
Ubuntu Trusty 14.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/dists/trusty-cm5.16.2.4505/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/ubuntu/trusty/amd64/cm/5.16.2.4505.tar.gz

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<code>usrepo.cloudera.com/p/cm5/ubuntu/trusty/amd64/cm/dists/trusty-cm5.16.2.4505/</code>	<code>usrepo.cloudera.com/p/cm5/ubuntu/trusty/amd64/cm/cloudera.list</code>	<code>usrepo.cloudera.com/p/cm5/cm/5/cloudera-trusty-5.16.2.4505.tar.gz</code>

Cloudera Manager 5.16.2

Release Date: June 18, 2019 **Status:** Production

Documentation:

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Table 22: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/5.16.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm/5/5.16.2/cm5.16.2.tar.gz
yum RHEL/CentOS/Oracle 6	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/5.16.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm/5/5.16.2/cm5.16.2.tar.gz
yum RHEL/CentOS/Oracle 5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/5.16.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm/5/5.16.2/cm5.16.2.tar.gz
apt Debian Jessie 8	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/dists/jessie-cm5.16.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm/5/5.16.2/cm5.16.2.tar.gz
apt Debian Wheezy 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/dists/wheezy-cm5.16.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm/5/5.16.2/cm5.16.2.tar.gz
zypper/YaST SLES 12	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/5.16.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/cloudera-cm.repo	https://s3.amazonaws.com/cloudera-repo/cm/5/5.16.2/cm5.16.2.tar.gz
zypper/YaST SLES 11	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/5.16.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm/5/5.16.2/cm5.16.2.tar.gz
Ubuntu Xenial 16.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/xenial/amd64/cm/dists/xenial-cm5.16.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/xenial/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm/5/5.16.2/cm5.16.2.tar.gz
Ubuntu Trusty 14.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/dists/trusty-cm5.16.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm/5/5.16.2/cm5.16.2.tar.gz

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<code>usrepo.cloudera.com/p/cm5/ubuntu/trusty/amd64/cm/dists/trusty-cm5.16.2/</code>	<code>usrepo.cloudera.com/p/cm5/ubuntu/trusty/amd64/cm/cloudera.list</code>	<code>usrepo.cloudera.com/p/cm5/cm/5/cloudera-trusty-5.16.2.tar.gz</code>

Previous Versions of Cloudera Manager 5



Important: The Cloudera Manager installer always installs the most recent version of Cloudera Manager 5. To install a previous version of Cloudera Manager 5, configure the repository for the release you want to install in [Step 1: Configure a Repository for Cloudera Manager](#).

- **5.16.1**



Note: The first release of Cloudera Manager 5.16 is numbered 5.16.1.

Release Date: November 28, 2018 **Status:** Production

Documentation:

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Table 23: Repositories

Type	Location	Repo file for installation of latest version of Cludera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 71	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ 5.16.1/	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.16.1.tar.gz
yum RHEL/CentOS/Oracle 6	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ 5.16.1/	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.16.1.tar.gz
yum RHEL/CentOS/Oracle 5	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ 5.16.1/	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.16.1.tar.gz
apt Debian Jessie 8	https:// s3.amazonaws.com/ p/cm5/debian/ jessie/amd64/cm/ dists/ jessie-cm5.16.1/	https:// s3.amazonaws.com/ p/cm5/debian/ jessie/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.16.1.tar.gz
apt Debian Wheezy 7	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.16.1/	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.16.1.tar.gz
zypper/YaST SLES 12	https:// s3.amazonaws.com/ p/cm5/sles/12/ x86_64/cm/ 5.16.1/	https:// s3.amazonaws.com/ p/cm5/sles/12/ x86_64/cm/ cloudera-cm.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.16.1.tar.gz
zypper/YaST SLES 11	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ 5.16.1/	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.16.1.tar.gz
Ubuntu Xenial 16.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ xenial/amd64/cm/ dists/ xenial-cm5.16.1/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ xenial/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.16.1.tar.gz
Ubuntu Trusty 14.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/	https:// s3.amazonaws.com/ p/cm5/ubuntu/	https:// s3.amazonaws.com/ p/cm5/cm/5/

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	trusty/amd64/cm/ dists/ trusty-cm5.16.1/	trusty/amd64/cm/ cloudera.list	cloudera-trusty-5.16.2.tar.gz

- **5.15.2**

Release Date: December 13, 2018 **Status:** Production

Documentation:

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Table 24: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/5.15.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-cm5.15.2.tar.gz
yum RHEL/CentOS/Oracle 6	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/5.15.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-cm5.15.2.tar.gz
yum RHEL/CentOS/Oracle 5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/5.15.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-cm5.15.2.tar.gz
apt Debian Jessie 8	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/dists/jessie-cm5.15.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-jss5.15.2.tar.gz
apt Debian Wheezy 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/dists/wheezy-cm5.15.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-wzy5.15.2.tar.gz
zypper/YaST SLES 12	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/5.15.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/cloudera-cm.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-sls5.15.2.tar.gz
zypper/YaST SLES 11	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/5.15.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-sls11.15.2.tar.gz
Ubuntu Xenial 16.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/xenial/amd64/cm/dists/xenial-cm5.15.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/xenial/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-xen16.04.tar.gz
Ubuntu Trusty 14.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<pre>trusty/amd64/cm/ dists/ trusty-cm5.15.2/</pre>	<pre>trusty/amd64/cm/ cloudera.list</pre>	<pre>cloudera-trusty-cm5.15.2.tar.gz</pre>
Ubuntu Precise 12.04	<pre>https:// s3.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.15.2/</pre>	<pre>https:// s3.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</pre>	<pre>https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-precise-cm5.15.2.tar.gz</pre>

- **5.15.1**

Release Date: August 23, 2018 **Status:** Production

Documentation:

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Table 25: Repositories

Type	Location	Repo file for installation of latest version of Cludera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ 5.15.1/	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.15.1.tar.gz
yum RHEL/CentOS/Oracle 6	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ 5.15.1/	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.15.1.tar.gz
yum RHEL/CentOS/Oracle 5	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ 5.15.1/	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.15.1.tar.gz
apt Debian Jessie 8	https:// s3.amazonaws.com/ p/cm5/debian/ jessie/amd64/cm/ dists/ jessie-cm5.15.1/	https:// s3.amazonaws.com/ p/cm5/debian/ jessie/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-jessie.15.1.tar.gz
apt Debian Wheezy 7	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.15.1/	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-wheezy.15.1.tar.gz
zypper/YaST SLES 12	https:// s3.amazonaws.com/ p/cm5/sles/12/ x86_64/cm/ 5.15.1/	https:// s3.amazonaws.com/ p/cm5/sles/12/ x86_64/cm/ cloudera-cm.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-sles.15.1.tar.gz
zypper/YaST SLES 11	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ 5.15.1/	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-sles.15.1.tar.gz
Ubuntu Xenial 16.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ xenial/amd64/cm/ dists/ xenial-cm5.15.1/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ xenial/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-xenial.15.1.tar.gz
Ubuntu Trusty 14.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/	https:// s3.amazonaws.com/ p/cm5/ubuntu/	https:// s3.amazonaws.com/ p/cm5/cm/5/

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<pre>trusty/amd64/cm/ dists/ trusty-cm5.15.1/</pre>	<pre>trusty/amd64/cm/ cloudera.list</pre>	<pre>cloudera-trusty-cm5.16.2.tar.gz</pre>
Ubuntu Precise 12.04	<pre>https:// userassets.cloudera.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.15.1/</pre>	<pre>https:// userassets.cloudera.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</pre>	<pre>https:// userassets.cloudera.com/ p/cm5/cm/5/ cloudera-precise-cm5.16.2.tar.gz</pre>

- **5.15.0**

Release Date: June 15, 2018 **Status:** Production

Documentation:

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<pre>trusty/amd64/cm/ dists/ trusty-cm5.15.0/</pre>	<pre>trusty/amd64/cm/ cloudera.list</pre>	<pre>cloudera-trusty-cm5.15.0.tar.gz</pre>
Ubuntu Precise 12.04	<pre>https:// userassets.cloudera.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.15.0/</pre>	<pre>https:// userassets.cloudera.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</pre>	<pre>https:// userassets.cloudera.com/ p/cm5/cm/5/ cloudera-precise-cm5.15.0.tar.gz</pre>

- **5.14.4**

Release Date: July 12, 2018 **Status:** Production

Documentation:

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Table 27: Repositories

Type	Location	Repo file for installation of latest version of Cludera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/5.14.4/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-cm5.14.86.tar.gz
yum RHEL/CentOS/Oracle 6	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/5.14.4/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-cm5.14.86.tar.gz
yum RHEL/CentOS/Oracle 5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/5.14.4/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-cm5.14.86.tar.gz
apt Debian Jessie 8	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/dists/jessie-cm5.14.4/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-jss5.14.86.tar.gz
apt Debian Wheezy 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/dists/wheezy-cm5.14.4/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-wzy5.14.86.tar.gz
zypper/YaST SLES 12	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/5.14.4/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/cloudera-cm.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-sls5.14.86.tar.gz
zypper/YaST SLES 11	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/5.14.4/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-sls5.14.86.tar.gz
Ubuntu Xenial 16.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/xenial/amd64/cm/dists/xenial-cm5.14.4/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/xenial/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-xen5.14.86.tar.gz
Ubuntu Trusty 14.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<pre>trusty/amd64/cm/ dists/ trusty-cm5.14.4/</pre>	<pre>trusty/amd64/cm/ cloudera.list</pre>	<pre>cloudera-trusty-5.14.4.tar.gz</pre>
Ubuntu Precise 12.04	<pre>https:// userassets.cloudera.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.14.4/</pre>	<pre>https:// userassets.cloudera.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</pre>	<pre>https:// userassets.cloudera.com/ p/cm5/cm/5/ cloudera-precise-5.14.4.tar.gz</pre>

- **5.13.3**

Release Date: April 4, 2018 **Status:** Production

Documentation:

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Table 28: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ 5.13.3/	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.13.3.tar.gz
yum RHEL/CentOS/Oracle 6	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ 5.13.3/	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.13.3.tar.gz
yum RHEL/CentOS/Oracle 5	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ 5.13.3/	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.13.3.tar.gz
apt Debian Jessie 8	https:// s3.amazonaws.com/ p/cm5/debian/ jessie/amd64/cm/ dists/ jessie-cm5.13.3/	https:// s3.amazonaws.com/ p/cm5/debian/ jessie/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.13.3.tar.gz
apt Debian Wheezy 7	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.13.3/	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.13.3.tar.gz
zypper/YaST SLES 12	https:// s3.amazonaws.com/ p/cm5/sles/12/ x86_64/cm/ 5.13.3/	https:// s3.amazonaws.com/ p/cm5/sles/12/ x86_64/cm/ cloudera-cm.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.13.3.tar.gz
zypper/YaST SLES 11	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ 5.13.3/	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.13.3.tar.gz
Ubuntu Xenial 16.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ xenial/amd64/cm/ dists/ xenial-cm5.13.3/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ xenial/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.13.3.tar.gz
Ubuntu Trusty 14.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/	https:// s3.amazonaws.com/ p/cm5/ubuntu/	https:// s3.amazonaws.com/ p/cm5/cm/5/

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<pre>trusty/amd64/cm/ dists/ trusty-cm5.13.3/</pre>	<pre>trusty/amd64/cm/ cloudera.list</pre>	<pre>cloudera-manager-5.13.3-amd64.tar.gz</pre>
Ubuntu Precise 12.04	<pre>https:// userassets.cloudera.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.13.3/</pre>	<pre>https:// userassets.cloudera.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</pre>	<pre>https:// userassets.cloudera.com/ p/cm5/cm/5/ cloudera-precise-5.13.3-amd64.tar.gz</pre>

- **5.12.2**

Release Date: December 15, 2017 **Status:** Production

Documentation:

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Table 29: Repositories

Type	Location	Repo file for installation of latest version of Cludera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/redhat/7/x86_64/cm/5.12.2/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/redhat/7/x86_64/cm/cludera-manager.repo	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/cm/5/cludera-cm5.12.2.tar.gz
yum RHEL/CentOS/Oracle 6	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/redhat/6/x86_64/cm/5.12.2/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/redhat/6/x86_64/cm/cludera-manager.repo	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/cm/5/cludera-cm5.12.2.tar.gz
yum RHEL/CentOS/Oracle 5	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/redhat/5/x86_64/cm/5.12.2/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/redhat/5/x86_64/cm/cludera-manager.repo	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/cm/5/cludera-cm5.12.2.tar.gz
apt Debian Jessie 8	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/debian/jessie/amd64/cm/dists/jessie-cm5.12.2/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/debian/jessie/amd64/cm/cludera.list	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/cm/5/cludera-jessie.12.2.tar.gz
apt Debian Wheezy 7	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/debian/wheezy/amd64/cm/dists/wheezy-cm5.12.2/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/debian/wheezy/amd64/cm/cludera.list	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/cm/5/cludera-wheezy.12.2.tar.gz
zypper/YaST SLES 12	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/sles/12/x86_64/cm/5.12.2/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/sles/12/x86_64/cm/cludera-cm.repo	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/cm/5/cludera-sles.12.2.tar.gz
zypper/YaST SLES 11	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/sles/11/x86_64/cm/5.12.2/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/sles/11/x86_64/cm/cludera-manager.repo	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/cm/5/cludera-sles.11.12.2.tar.gz
Ubuntu Xenial 16.04	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/ubuntu/xenial/amd64/cm/dists/xenial-cm5.12.2/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/ubuntu/xenial/amd64/cm/cludera.list	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/cm/5/cludera-xenial.12.2.tar.gz
Ubuntu Trusty 14.04	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/ubuntu/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/ubuntu/	https://s3.amazonaws.com/cludera-repo-us-east-1-us-west-2/p/cm5/cm/5/

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<pre>trusty/amd64/cm/ dists/ trusty-cm5.12.2/</pre>	<pre>trusty/amd64/cm/ cloudera.list</pre>	<pre>cloudera-trusty-5.12.2.tar.gz</pre>
Ubuntu Precise 12.04	<pre>https:// sraa7asv0d1ie2ubac/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.12.2/</pre>	<pre>https:// sraa7asv0d1ie2ubac/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</pre>	<pre>https:// sraa7asv0d1ie2ubac/ p/cm5/cm/5/ cloudera-precise-5.12.2.tar.gz</pre>

- **5.11.2**

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation](#)
- [Cloudera Upgrade](#)

Table 30: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/5.11.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-cm5.11.2.tar.gz
yum RHEL/CentOS/Oracle 6	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/5.11.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-cm5.11.2.tar.gz
yum RHEL/CentOS/Oracle 5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/5.11.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-cm5.11.2.tar.gz
apt Debian Jessie 8	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/dists/jessie-cm5.11.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-jss5.11.2.tar.gz
apt Debian Wheezy 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/dists/wheezy-cm5.11.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-wzy5.11.2.tar.gz
zypper/YaST SLES 12	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/5.11.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/cloudera-cm.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-sls5.11.2.tar.gz
zypper/YaST SLES 11	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/5.11.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-sls11.2.tar.gz
Ubuntu Trusty 14.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/dists/trusty-cm5.11.2/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-usty5.11.2.tar.gz
Ubuntu Precise 12.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<code>precise/amd64/ cm/dists/ precise-cm5.11.2/</code>	<code>precise/amd64/ cm/cloudera.list</code>	<code>cloudera-precis5.11.2.tar.gz</code>

- **5.10.2**
 - [Cloudera Manager 5 Release Notes](#) on page 2472
 - [Cloudera Installation](#)
 - [Cloudera Upgrade](#)

Table 31: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ 5.10.2/	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.10.2.tar.gz
yum RHEL/CentOS/Oracle 6	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ 5.10.2/ ../topics/ cloudera-cm5.10.2.tar.gz	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.10.2.tar.gz
yum RHEL/CentOS/Oracle 5	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ 5.10.2/	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.10.2.tar.gz
apt Debian Jessie 8	https:// s3.amazonaws.com/ p/cm5/debian/ jessie/amd64/cm/ dists/ jessie-cm5.10.2/	https:// s3.amazonaws.com/ p/cm5/debian/ jessie/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-jessie.10.2.tar.gz
apt Debian Wheezy 7	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.10.2/	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-wheezy.10.2.tar.gz
zypper/YaST SLES 12	https:// s3.amazonaws.com/ p/cm5/sles/12/ x86_64/cm/ 5.10.2/	https:// s3.amazonaws.com/ p/cm5/sles/12/ x86_64/cm/ cloudera-cm.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-sles.10.2.tar.gz
zypper/YaST SLES 11	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ 5.10.2/	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-sles.10.2.tar.gz
Ubuntu Trusty 14.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ trusty/amd64/cm/ dists/ trusty-cm5.10.2/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ trusty/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-trusty.10.2.tar.gz

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
Ubuntu Precise 12.04	<code>https:// us-east-1.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.10.2/</code>	<code>https:// us-east-1.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</code>	<code>https:// us-east-1.amazonaws.com/ p/cm5/cm/5/ cloudera-precise-12.04.tar.gz</code>

- 5.9.3
 - [Cloudera Manager 5 Release Notes](#) on page 2472
 - [Cloudera Installation](#)
 - [Cloudera Upgrade](#)

Table 32: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/5.9.3/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-cm5-5.9.3.tar.gz
yum RHEL/CentOS/Oracle 6	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/5.9.3/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-cm5-5.9.3.tar.gz
yum RHEL/CentOS/Oracle 5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/5.9.3/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-cm5-5.9.3.tar.gz
apt Debian Jessie 8	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/dists/jessie-cm5.9.3/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-jessie-5.9.3.tar.gz
apt Debian Wheezy 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/dists/wheezy-cm5.9.3/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-wheezy-5.9.3.tar.gz
zypper/YaST SLES 12	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/5.9.3/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/12/x86_64/cm/cloudera-cm.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-sles-5.9.3.tar.gz
zypper/YaST SLES 11	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/5.9.3/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-sles-5.9.3.tar.gz
Ubuntu Trusty 14.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/dists/trusty-cm5.9.3/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-trusty-5.9.3.tar.gz
Ubuntu Precise 12.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<code>precise/amd64/ cm/dists/ precise-cm5.9.3/</code>	<code>precise/amd64/ cm/cloudera.list</code>	<code>cloudera-manager-5.9.3-0414.gz</code>

- **5.8.5**
 - [Cloudera Manager 5 Release Notes](#) on page 2472
 - [Cloudera Installation](#)
 - [Cloudera Upgrade](#)

Table 33: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/5.8.5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/7/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-cm5-5.8.5.tar.gz
yum RHEL/CentOS/Oracle 6	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/5.8.5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-cm5-5.8.5.tar.gz
yum RHEL/CentOS/Oracle 5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/5.8.5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-cm5-5.8.5.tar.gz
apt Debian Jessie 8	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/dists/jessie-cm5.8.5/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/jessie/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-jessie-5.8.5.tar.gz
apt Debian Wheezy 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/dists/wheezy-cm5.8.5/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-wheezy-5.8.5.tar.gz
zypper/YaST SLES 11	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/5.8.5/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-sles-5.8.5.tar.gz
Ubuntu Trusty 14.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/dists/trusty-cm5.8.5/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-trusty-5.8.5.tar.gz
Ubuntu Precise 12.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/precise/amd64/cm/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/precise/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/p/cm5/cm/5/cloudera-precise-5.8.5.tar.gz

- 5.7.6
 - [Cloudera Manager 5 Release Notes](#) on page 2472

Version and Download Information

- [Cloudera Installation and Upgrade](#)

Table 34: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/5.7.5/	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.7.5.tar.gz
yum RHEL/CentOS/Oracle 6	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/5.7.5/	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.7.5.tar.gz
yum RHEL/CentOS/Oracle 5	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/5.7.5/	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.7.5.tar.gz
apt Debian Wheezy 7	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.7.5/	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.7.5.tar.gz
apt Debian Squeeze 6	https:// s3.amazonaws.com/ p/cm5/debian/ squeeze/amd64/ cm/dists/ squeeze-cm5.7/	https:// s3.amazonaws.com/ p/cm5/debian/ squeeze/amd64/ cm/cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.7.5.tar.gz
zypper/YaST SLES 11	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/5.7/	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.7.5.tar.gz
Ubuntu Trusty 14.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ trusty/amd64/cm/ dists/ trusty-cm5.7.5/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ trusty/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.7.5.tar.gz
Ubuntu Precise 12.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.7.5/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.7.5.tar.gz
Ubuntu Lucid 10.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ lucid/amd64/cm/ dists/ lucid-cm5.7.5/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ lucid/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.7.5.tar.gz

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	p/cm5/ubuntu/ lucid/amd64/cm/ dists/ lucid-cm5.7/	p/cm5/ubuntu/ lucid/amd64/cm/ cloudera.list	p/cm5/cm/5/ cloudera-lucid-5.7.tar.gz

- 5.6.1
 - [Cloudera Manager 5 Release Notes](#) on page 2472
 - [Cloudera Installation and Upgrade](#)

Table 35: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/5.6.1/	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.6.1.tar.gz
yum RHEL/CentOS/Oracle 6	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/5.6.1/	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.6.1.tar.gz
yum RHEL/CentOS/Oracle 5	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/5.6.1/	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.6.1.tar.gz
apt Debian Wheezy 7	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.6.1/	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.6.1.tar.gz
apt Debian Squeeze 6	https:// s3.amazonaws.com/ p/cm5/debian/ squeeze/amd64/ cm/dists/ squeeze-cm5.6.1/	https:// s3.amazonaws.com/ p/cm5/debian/ squeeze/amd64/ cm/cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.6.1.tar.gz
zypper/YaST SLES 11	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/5.6.1/	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.6.1.tar.gz
Ubuntu Trusty 14.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ trusty/amd64/cm/ dists/ trusty-cm5.6.1/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ trusty/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.6.1.tar.gz
Ubuntu Precise 12.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.6.1/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.6.1.tar.gz
Ubuntu Lucid 10.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ lucid/amd64/cm/ dists/ lucid-cm5.6.1/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ lucid/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5.6.1.tar.gz

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<code>p/cm5/ubuntu/ lucid/amd64/cm/ dists/ lucid-cm5.6.1/</code>	<code>p/cm5/ubuntu/ lucid/amd64/cm/ cloudera.list</code>	<code>p/cm5/cm/5/ cloudera-lcdm561.tar.gz</code>

- **5.5.6**
 - [Cloudera Manager 5 Release Notes](#) on page 2472
 - [Cloudera Installation and Upgrade](#)

Table 36: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 7	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/5.5.6/	https:// s3.amazonaws.com/ p/cm5/redhat/7/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.5.6.tar.gz
yum RHEL/CentOS/Oracle 6	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/5.5.6/	https:// s3.amazonaws.com/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.5.6.tar.gz
yum RHEL/CentOS/Oracle 5	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/5.5.6/	https:// s3.amazonaws.com/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.5.6.tar.gz
apt Debian Wheezy 7	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.5.6/	https:// s3.amazonaws.com/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.5.6.tar.gz
apt Debian Squeeze 6	https:// s3.amazonaws.com/ p/cm5/debian/ squeeze/amd64/ cm/dists/ squeeze-cm5.5.6/	https:// s3.amazonaws.com/ p/cm5/debian/ squeeze/amd64/ cm/cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.5.6.tar.gz
zypper/YaST SLES 11	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/5.5.6/	https:// s3.amazonaws.com/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.5.6.tar.gz
Ubuntu Trusty 14.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ trusty/amd64/cm/ dists/ trusty-cm5.5.6/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ trusty/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.5.6.tar.gz
Ubuntu Precise 12.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.5.6/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.5.6.tar.gz
Ubuntu Lucid 10.04	https:// s3.amazonaws.com/ p/cm5/ubuntu/ lucid/amd64/cm/ dists/ lucid-cm5.5.6/	https:// s3.amazonaws.com/ p/cm5/ubuntu/ lucid/amd64/cm/ cloudera.list	https:// s3.amazonaws.com/ p/cm5/cm/5/ cloudera-cm5-5.5.6.tar.gz

Version and Download Information

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
	<code>p/cm5/ubuntu/ lucid/amd64/cm/ dists/ lucid-cm5.5.6/</code>	<code>p/cm5/ubuntu/ lucid/amd64/cm/ cloudera.list</code>	<code>p/cm5/cm/5/ cloudera-lcdm56.tar.gz</code>

- **5.4.10**
 - [Cloudera Manager 5 Release Notes](#) on page 2472
 - [Cloudera Installation and Upgrade](#)

Table 37: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 6	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/5.4.10/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/6/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-cm54086.tar.gz
yum RHEL/CentOS/Oracle 5	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/5.4.10/	https://s3.amazonaws.com/cloudera-repo/p/cm5/redhat/5/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-cm54086.tar.gz
apt Debian Wheezy 7	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/dists/wheezy-cm5.4.10/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/wheezy/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-wzy54086.tar.gz
apt Debian Squeeze 6	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/squeeze/amd64/cm/dists/squeeze-cm5.4.10/	https://s3.amazonaws.com/cloudera-repo/p/cm5/debian/squeeze/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-sze54086.tar.gz
zypper/YaST SLES 11	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/5.4.10/	https://s3.amazonaws.com/cloudera-repo/p/cm5/sles/11/x86_64/cm/cloudera-manager.repo	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-sls1154086.tar.gz
Ubuntu Trusty 14.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/dists/trusty-cm5.4.10/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/trusty/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-tusy14086.tar.gz
Ubuntu Precise 12.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/precise/amd64/cm/dists/precise-cm5.4.10/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/precise/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-prec12086.tar.gz
Ubuntu Lucid 10.04	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/lucid/amd64/cm/dists/lucid-cm5.4.10/	https://s3.amazonaws.com/cloudera-repo/p/cm5/ubuntu/lucid/amd64/cm/cloudera.list	https://s3.amazonaws.com/cloudera-repo/cm5/cm/5/cloudera-luc1054086.tar.gz

- 5.3.10

Version and Download Information

- [Cloudera Manager 5 Release Notes](#) on page 2472
- [Cloudera Installation and Upgrade](#)

Table 38: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 6	<code>https:// usm1.safelinkcloud.com/ p/cm5/redhat/6/ x86_64/cm/ 5.3.10/</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/cm/5/ cloudera-cm53086.tar.gz</code>
yum RHEL/CentOS/Oracle 5	<code>https:// usm1.safelinkcloud.com/ p/cm5/redhat/5/ x86_64/cm/ 5.3.10/</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/cm/5/ cloudera-cm53086.tar.gz</code>
apt Debian Wheezy 7	<code>https:// usm1.safelinkcloud.com/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.3.10/</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/cm/5/ cloudera-wzy-cm53086.tar.gz</code>
apt Debian Squeeze 6	<code>https:// usm1.safelinkcloud.com/ p/cm5/debian/ squeeze/amd64/ cm/dists/ squeeze-cm5.3.10/</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/debian/ squeeze/amd64/ cm/cloudera.list</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/cm/5/ cloudera-sze-cm53086.tar.gz</code>
zypper/YaST SLES 11	<code>https:// usm1.safelinkcloud.com/ p/cm5/sles/11/ x86_64/cm/ 5.3.10/</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/cm/5/ cloudera-sls-cm53086.tar.gz</code>
Ubuntu Trusty 14.04	<code>https:// usm1.safelinkcloud.com/ p/cm5/ubuntu/ trusty/amd64/cm/ 5.3.10/</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/ubuntu/ trusty/amd64/cm/ cloudera.list</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/cm/5/ cloudera-trusy-cm53086.tar.gz</code>
Ubuntu Precise 12.04	<code>https:// usm1.safelinkcloud.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.3.10/</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/cm/5/ cloudera-precis-cm53086.tar.gz</code>
Ubuntu Lucid 10.04	<code>https:// usm1.safelinkcloud.com/ p/cm5/ubuntu/ lucid/amd64/cm/ dists/ lucid-cm5.3.10/</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/ubuntu/ lucid/amd64/cm/ cloudera.list</code>	<code>https:// usm1.safelinkcloud.com/ p/cm5/cm/5/ cloudera-lucid-cm53086.tar.gz</code>

- 5.2.7

- [Cloudera Manager 5 Release Notes](#) on page 2472

Version and Download Information

- [Cloudera Installation and Upgrade](#)

- [Cloudera Manager Release Notes](#)
- [Cloudera Manager Installation Guide](#)

Table 40: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 6	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/redhat/6/ x86_64/cm/5.1.6/</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/cm/5/ cloudera-5.16.2.tar.gz</code>
yum RHEL/CentOS/Oracle 5	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/redhat/5/ x86_64/cm/5.1.6/</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/cm/5/ cloudera-5.16.2.tar.gz</code>
apt Debian Wheezy 7	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.1.6/</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/cm/5/ cloudera-wheezy-5.16.2.tar.gz</code>
apt Debian Squeeze 6	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/debian/ squeeze/amd64/ cm/dists/ squeeze-cm5.1.6/</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/debian/ squeeze/amd64/ cm/cloudera.list</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/cm/5/ cloudera-squeeze-5.16.2.tar.gz</code>
zypper/YaST SLES 11	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/sles/11/ x86_64/cm/5.1.6/</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/cm/5/ cloudera-sles-5.16.2.tar.gz</code>
Ubuntu Precise 12.04	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.1.6/</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/cm/5/ cloudera-precise-5.16.2.tar.gz</code>
Ubuntu Lucid 10.04	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/ubuntu/ lucid/amd64/cm/ dists/ lucid-cm5.1.6/</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/ubuntu/ lucid/amd64/cm/ cloudera.list</code>	<code>https:// usm1tasvodad1ie.duba.com/ p/cm5/cm/5/ cloudera-lucid-5.16.2.tar.gz</code>

- **5.0.7**
 - [Cloudera Manager Release Notes](#)
 - [Cloudera Manager Installation Guide](#)

Table 41: Repositories

Type	Location	Repo file for installation of latest version of Cloudera Manager 5 (5.16.2)	Tarball File
yum RHEL/CentOS/Oracle 6	<code>https:// usartasvodadiedubam/ p/cm5/redhat/6/ x86_64/cm/5.0.7/</code>	<code>https:// usartasvodadiedubam/ p/cm5/redhat/6/ x86_64/cm/ cloudera-manager.repo</code>	<code>https:// usartasvodadiedubam/ p/cm5/cm/5/ cloudera-cm5.0.7.86.tar.gz</code>
yum RHEL/CentOS/Oracle 5	<code>https:// usartasvodadiedubam/ p/cm5/redhat/5/ x86_64/cm/5.0.7/</code>	<code>https:// usartasvodadiedubam/ p/cm5/redhat/5/ x86_64/cm/ cloudera-manager.repo</code>	<code>https:// usartasvodadiedubam/ p/cm5/cm/5/ cloudera-cm5.0.7.86.tar.gz</code>
apt Debian Wheezy 7	<code>https:// usartasvodadiedubam/ p/cm5/debian/ wheezy/amd64/cm/ dists/ wheezy-cm5.0.7/</code>	<code>https:// usartasvodadiedubam/ p/cm5/debian/ wheezy/amd64/cm/ cloudera.list</code>	<code>https:// usartasvodadiedubam/ p/cm5/cm/5/ cloudera-wzy-cm5.0.7.86.tar.gz</code>
apt Debian Squeeze 6	<code>https:// usartasvodadiedubam/ p/cm5/debian/ squeeze/amd64/ cm/dists/ squeeze-cm5.0.7/</code>	<code>https:// usartasvodadiedubam/ p/cm5/debian/ squeeze/amd64/ cm/cloudera.list</code>	<code>https:// usartasvodadiedubam/ p/cm5/cm/5/ cloudera-sze-cm5.0.7.86.tar.gz</code>
zypper/YaST SLES 11	<code>https:// usartasvodadiedubam/ p/cm5/sles/11/ x86_64/cm/5.0.7/</code>	<code>https:// usartasvodadiedubam/ p/cm5/sles/11/ x86_64/cm/ cloudera-manager.repo</code>	<code>https:// usartasvodadiedubam/ p/cm5/cm/5/ cloudera-sls-cm5.0.7.86.tar.gz</code>
Ubuntu Precise 12.04	<code>https:// usartasvodadiedubam/ p/cm5/ubuntu/ precise/amd64/ cm/dists/ precise-cm5.0.7/</code>	<code>https:// usartasvodadiedubam/ p/cm5/ubuntu/ precise/amd64/ cm/cloudera.list</code>	<code>https:// usartasvodadiedubam/ p/cm5/cm/5/ cloudera-prc-cm5.0.7.86.tar.gz</code>
Ubuntu Lucid 10.04	<code>https:// usartasvodadiedubam/ p/cm5/ubuntu/ lucid/amd64/cm/ dists/ lucid-cm5.0.7/</code>	<code>https:// usartasvodadiedubam/ p/cm5/ubuntu/ lucid/amd64/cm/ cloudera.list</code>	<code>https:// usartasvodadiedubam/ p/cm5/cm/5/ cloudera-luc-cm5.0.7.86.tar.gz</code>

CDH Version and Packaging Information

For installation instructions for CDH 5, see [Cloudera Installation](#).

If you are using a platform that is not supported by Debian, RPM, or SUSE packages, you can download the patched sources – see the Patched Tarball links for the version you want.

Version and Download Information

CDH is available in the following releases:

- CDH 5 is the current stable release, based on Apache Hadoop 2.3.0 or later.

For information on the exact Apache Hadoop version included in each CDH version, see [CDH 5 Packaging and Tarball Information](#) on page 230. For platforms not supported by Debian, RPM, or SLES packages, download the tarball with the appropriate patch-level.



Important:

- Cloudera Manager 4 and CDH 4 reached End of Maintenance (EOM) on August 9, 2015.
- CDH 3 reached End of Maintenance (EOM) on June 20, 2013.

- CDH 4 is at End of Maintenance. No further updates are planned for this release.
- CDH 3 is at End of Maintenance. No further updates are planned for this release.
- CDH 2 is at End of Maintenance. No further updates are planned for this release.
- CDH 1 is at End of Maintenance. No further updates are planned for this release.



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH Download Information

These topics describe download information for CDH 5.



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.16.x Download Information



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

The section describes download information for CDH 5.16.x releases. For a different release, see .

CDH 5.16.2

Release Date: June 18, 2019 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.16.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.16.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.16.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.16.2/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.16.2/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.16.1

Release Date: November 28, 2018 **Status:** Production

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.16.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.16.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.16.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.16.1/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.16.1/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.15.x Download Information

The section describes download information for CDH 5.15.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.15.2

Release Date: December 13, 2018 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.15.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.15.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.15.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.15.2/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.15.2/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.15.2/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.15.1

Release Date: August 23, 2018 **Status:** Production

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.15.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.15.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.15.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.15.1/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.15.1/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.15.1/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.15.0

Release Date: May 3, 2018 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.15.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.15.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.15.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.15.0/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.15.0/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.15.0/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.14.x Download Information

The section describes download information for CDH 5.14.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.14.4

Release Date: July 12, 2018 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.14.4/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.14.4/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.14.4/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.14.4/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.14.4/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.14.4/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.14.2

Release Date: April 11, 2018 **Status:** Production

Repository Type	Location	Repo file

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.14.2/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.14.2/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.14.2/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Debian	<p>Jessie: https://username:password@archive.cloudera.com/p/cdh5/debian/jessie/amd64/cdh/dists/jessie-cdh5.14.2/</p> <p>Wheezy: https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/</p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh/dists/precise-cdh5.14.2/</p> <p>Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/dists/trusty-cdh5.14.2/</p>	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.14.0

Release Date: January 26, 2018 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.14.0/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.14.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.14.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.14.0/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.14.0/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.14.0/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.13.x Download Information

The section describes download information for CDH 5.13.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

Version and Download Information

CDH 5.13.3

Release Date: April 4, 2018 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.13.3/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.13.3/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.13.3/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.13.3/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.13.3/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.13.3/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.13.2

Release Date: February 23, 2018 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.13.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.13.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.13.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.13.2/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.13.2/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.13.2/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.13.1

Release Date: December 13, 2017 **Status:** Production

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.13.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.13.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.13.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.13.1/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.13.1/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.13.1/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.13.0

Release Date: October 12, 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.13.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.13.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.13.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.13.0/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.13.0/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.13.0/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.12.x Download Information

The section describes download information for CDH 5.12.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.12.2

Release Date: December 15, 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.12.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.12.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.12.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.12.2/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.12.2/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.12.2/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.12.1

Release Date: September 6, 2017 **Status:** Production

Repository Type	Location	Repo file

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.12.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.12.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.12.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.12.1/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.12.1/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.12.1/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.12.0

Release Date: July 13, 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.12.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.12.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.12.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.12.0/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.12.0/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.12.0/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.11.x Download Information

The section describes download information for CDH 5.11.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

Version and Download Information

CDH 5.11.2

Release Date: August 24, 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.11.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.11.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.11.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.11.2/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.11.2/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.11.2/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.11.1

Release Date: June 15, 2017 **Status:** Production

Repository Type	Location	Repo file

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.11.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.11.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.11.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.11.1/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.11.1/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.11.1/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.11.0

Release Date: April 18, 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.11.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.11.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.11.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.11.0/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.11.0/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.11.0/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.10.x Download Information

The section describes download information for CDH 5.10.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.10.2

Release Date: June 30, 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.10.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.10.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.10.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.10.2/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.10.2/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.10.2/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.10.1

Release Date: March 31, 2017 **Status:** Production

Repository Type	Location	Repo file

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.10.1/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.10.1/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.10.1/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Debian	<p>Jessie: https://username:password@archive.cloudera.com/p/cdh5/debian/jessie/amd64/cdh/dists/jessie-cdh5.10.1/</p> <p>Wheezy: https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/</p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh/dists/precise-cdh5.10.1/</p> <p>Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/dists/trusty-cdh5.10.1/</p>	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.10.0

Release Date: January 31, 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.10.0/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.10.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.10.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.10.0/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.10.0/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.10.0/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.9.x Download Information

The section describes download information for CDH 5.9.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

Version and Download Information

CDH 5.9.3

Release Date: July 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.9.3/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.9.3/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.9.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.9.3/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/dists/ wheezy-cdh5.9.3/</code></p>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/12/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.9.3/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.9.3/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.9.2

Release Date: April 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.9.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.9.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.9.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.9.2/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/dists/ wheezy-cdh5.9.2/</code></p>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/12/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.9.2/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.9.2/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.9.1

Release Date: December 2016 **Status:** Production

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.9.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.9.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.9.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.9.1/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/dists/ wheezy-cdh5.9.1/</code></p>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/12/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.9.1/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.9.1/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.9.0

Release Date: October 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.9.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.9.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.9.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/dists/ jessie-cdh5.9.0/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)		<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/dists/ precise-cdh5.9.0/</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/dists/ trusty-cdh5.9.0/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.8.x Download Information

The section describes download information for CDH 5.8.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.8.5

Release Date: May 2017 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.8.5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.8.5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.8.5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Debian	<p>Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/</code></p> <p>Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code></p>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	<p>Precise:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code></p> <p>Trusty:</p> <p><code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</code></p>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.8.4

Release Date: February 2017 Status: Production

Repository Type	Location	Repo file

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.8.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.8.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.8.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Debian	<p>Jessie: https://username:password@archive.cloudera.com/p/cdh5/debian/jessie/amd64/cdh/</p> <p>Wheezy: https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/</p>	N/A
Zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	<p>Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh</p> <p>Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/</p>	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.8.3

Release Date: October 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.8.3/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.8.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.8.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Debian	<p>Jessie: https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/</p> <p>Wheezy: https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</p>	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	<p>Precise:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</p> <p>Trusty:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</p>	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.8.2

Release Date: September 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.8.2/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.8.2/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.8.2/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Debian	<p>Jessie: https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/</p> <p>Wheezy: https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</p>	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	<p>Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</p> <p>Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</p>	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.8.0

Release Date: July 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.8.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.8.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.8.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo

Repository Type	Location	Repo file
Debian	Jessie: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/jessie/ amd64/cdh/</code> Wheezy: <code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code> Trusty: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.7.x Download Information

The section describes download information for CDH 5.7.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

Version and Download Information

CDH 5.7.6

Release Date: February 2017 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.7.6/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.7.6/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.7.6/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
Zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.7.5

Release Date: November 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.7.5/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.7.5/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.7.5/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
Zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.7.4

Release Date: October 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.7.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.7.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.7.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo

Version and Download Information

Repository Type	Location	Repo file
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	<p>Precise:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</p> <p>Trusty:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</p>	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.7.3

Release Date: September 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.7.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.7.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.7.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/

Repository Type	Location	Repo file
	cdh/5/	cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.7.2

Release Date: July 2016 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.7.2/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.7.2/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.7.2/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
Zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh	N/A

Repository Type	Location	Repo file
	Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.7.1

Release Date: June 2016 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.7.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.7.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.7.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/	N/A

Repository Type	Location	Repo file
	p/cdh5/cdh/5/	

CDH 5.7.0

Release Date: April 2015 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.7.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.7.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.7.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.6.x Download Information

The section describes download information for CDH 5.6.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.6.1

Release Date: May 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.6.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.6.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.6.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/ Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.6.0

Release Date: November 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.6.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.6.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.6.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code> Trusty: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.5.x Download Information

The section describes download information for CDH 5.5.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.5.6

Release Date: December 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.5.6/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.5.6/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.5.6/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.5.5

Release Date: October 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.5.5/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.5.5/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.5.5/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
Zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	<p>Precise:</p> <p>https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh</p> <p>Trusty:</p> <p>https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/</p>	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.5.4

Release Date: May 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.5.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.5.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.5.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo

Repository Type	Location	Repo file
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.5.2

Release Date: February 2016 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.5.2/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.5.2/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.5.2/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/

Version and Download Information

Repository Type	Location	Repo file
	cdh/5/	cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.5.1

Release Date: December 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/5.5.1/	https://username:password@archive.cloudera.com/p/cdh5/redhat/7/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.5.1/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.5.1/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
Zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh	N/A

Repository Type	Location	Repo file
	Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.5.0

Release Date: November 2015 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 7	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/5.5.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/7/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.5.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.5.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
Zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/	N/A

Version and Download Information

Repository Type	Location	Repo file
	p/cdh5/cdh/5/	

CDH 5.4.x Download Information

The section describes download information for CDH 5.4.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.4.11

Release Date: October 2016 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.4.11/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.4.11/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.4.10

Release Date: April 2016 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.4.10/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.4.10/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo

Version and Download Information

Repository Type	Location	Repo file
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.4.9

Release Date: December 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.4.9/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.4.9/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh/	N/A

Repository Type	Location	Repo file
	<p>p/cdh5/ubuntu/precise/ amd64/cdh</p> <p>Trusty:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</p>	
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.4.8

Release Date: October 2015 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.4.8/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.4.8/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	<p>Precise:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</p> <p>Trusty:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</p>	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

Version and Download Information

CDH 5.4.7

Release Date: September 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.4.7/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.4.7/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code> Trusty: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.4.5

Release Date: August 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.4.5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.4.5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/</code>	N/A

Version and Download Information

Repository Type	Location	Repo file
	amd64/cdh/	
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.4.4

Release Date: July 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.4.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.4.4/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh	N/A

Repository Type	Location	Repo file
	Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.4.3

Release Date: June 2015 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.4.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.4.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.4.2

Release Date: May 2015 Status: Production

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.4.2/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.4.2/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.4.1

Release Date: May 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.4.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.4.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/	N/A

Repository Type	Location	Repo file
	amd64/cdh/	
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.4.0

Release Date: April 2015 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.4.0/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.4.0/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh	N/A

Version and Download Information

Repository Type	Location	Repo file
	Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.3.x Download Information

The section describes download information for CDH 5.3.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.3.10

Release Date: April 2016 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.3.10/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.3.10/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.3.9

Release Date: January 2016 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.3.9/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.3.9/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo

Version and Download Information

Repository Type	Location	Repo file
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh Trusty: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.3.8

Release Date: October 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.3.8/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.3.8/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh/	N/A

Repository Type	Location	Repo file
	<p>p/cdh5/ubuntu/precise/ amd64/cdh</p> <p>Trusty:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</p>	
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.3.6

Release Date: August 2015 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.3.6/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.3.6/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	<p>Precise:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</p> <p>Trusty:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</p>	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

Version and Download Information

CDH 5.3.5

Release Date: June 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.3.5/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.3.5/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.3.4

Release Date: June 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.3.4/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.3.4/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo

Repository Type	Location	Repo file
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.3.3

Release Date: April 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.3.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.3.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh/	N/A

Version and Download Information

Repository Type	Location	Repo file
	<p>p/cdh5/ubuntu/precise/amd64/cdh</p> <p>Trusty:</p> <p>https:// username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/</p>	
Patched Tarball	<p>https:// username:password@archive.cloudera.com/p/cdh5/cdh/5/</p>	N/A

CDH 5.3.2

Release Date: February 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<p>https:// username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.3.2/</p>	<p>https:// username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo</p>
Yum RHEL/CentOS/Oracle 5	<p>https:// username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.3.2/</p>	<p>https:// username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo</p>
Apt (Debian)	<p>https:// username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/</p>	N/A
zypper/YaST (SLES)	<p>https:// username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/</p>	<p>https:// username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo</p>
Ubuntu	<p>Precise:</p> <p>https:// username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh/</p> <p>Trusty:</p> <p>https:// username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/</p>	N/A
Patched Tarball	<p>https:// username:password@archive.cloudera.com/p/cdh5/cdh/5/</p>	N/A

CDH 5.3.1

Release Date: February 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.3.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.3.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.3.0

Release Date: December 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.3.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.3.0/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo

Version and Download Information

Repository Type	Location	Repo file
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code> Trusty: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.2.x Download Information

The section describes download information for CDH 5.2.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.2.6

Release Date: July 2015 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.2.6/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.2.6/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.2.5

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.2.5/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.2.5/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https://	N/A

Version and Download Information

Repository Type	Location	Repo file
	<code>username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/</code>	
zypper/YaST (SLES)	<code>https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/</code>	<code>https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh</code> Trusty: <code>https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/</code>	N/A
Patched Tarball	<code>https://username:password@archive.cloudera.com/p/cdh5/cdh/5/</code>	N/A

CDH 5.2.4

Release Date: February 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.2.4/</code>	<code>https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.2.4/</code>	<code>https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/</code>	N/A
zypper/YaST (SLES)	<code>https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/</code>	<code>https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/</code>	N/A

Repository Type	Location	Repo file
	amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.2.3



Note: There is no CDH 5.0.7, 5.1.1, 5.1.6, 5.2.2, 5.2.7, 5.3.7, or 5.4.6 release.

Version and Download Information

Release Date: February 2015 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.2.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.2.3/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	<p>Precise:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</p> <p>Trusty:</p> <p>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/</p>	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.2.1

Release Date: December 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.2.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.2.1/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https://	N/A

Repository Type	Location	Repo file
	<code>username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/</code>	
zypper/YaST (SLES)	<code>https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/</code>	<code>https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh</code> Trusty: <code>https://username:password@archive.cloudera.com/p/cdh5/ubuntu/trusty/amd64/cdh/</code>	N/A
Patched Tarball	<code>https://username:password@archive.cloudera.com/p/cdh5/cdh/5/</code>	N/A

CDH 5.2.0

Release Date: October 2014 Status: Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.2.0/</code>	<code>https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.2.0/</code>	<code>https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/</code>	N/A
zypper/YaST (SLES)	<code>https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/</code>	<code>https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/</code>	N/A

Version and Download Information

Repository Type	Location	Repo file
	amd64/cdh Trusty: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/trusty/ amd64/cdh/	
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.1.x Download Information

The section describes download information for CDH 5.1.x releases. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

CDH 5.1.5

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.1.5/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.1.5/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo
Apt (Debian)	https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/	N/A
zypper/YaST (SLES)	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/	https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.1.4

Release Date: December 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.1.4/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.1.4/	https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo

Version and Download Information

Repository Type	Location	Repo file
	<code>p/cdh5/redhat/5/x86_64/cdh/5.1.4/</code>	<code>p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.1.3

Release Date: September 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.1.3/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.1.3/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code>	N/A

Repository Type	Location	Repo file
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.1.2

**Note:**

There is no CDH 5.1.1 release. This skip in the CDH 5.x sequence allows the CDH and Cloudera Manager components of Cloudera Enterprise 5.1.2 to have consistent numbering.

Version and Download Information

Release Date: August 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.1.2/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.1.2/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https://username:password@archive.cloudera.com/p/cdh5/ubuntu/precise/amd64/cdh	N/A
Patched Tarball	https://username:password@archive.cloudera.com/p/cdh5/cdh/5/	N/A

CDH 5.1.0

Release Date: July 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/5.1.0/	https://username:password@archive.cloudera.com/p/cdh5/redhat/6/x86_64/cdh/cloudera-cdh5.repo
Yum RHEL/CentOS/Oracle 5	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/5.1.0/	https://username:password@archive.cloudera.com/p/cdh5/redhat/5/x86_64/cdh/cloudera-cdh5.repo
Apt (Debian)	https://username:password@archive.cloudera.com/p/cdh5/debian/wheezy/amd64/cdh/	N/A
zypper/YaST (SLES)	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/5.1.0/	https://username:password@archive.cloudera.com/p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo

Repository Type	Location	Repo file
	p/cdh5/sles/11/x86_64/cdh/5/	p/cdh5/sles/11/x86_64/cdh/cloudera-cdh5.repo
Ubuntu	Precise: https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh	N/A
Patched Tarball	https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/	N/A

CDH 5.0.x Download Information

The section describes download information for CDH 5.0.x. For a different release, see .



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

Version and Download Information

CDH 5.0.6

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.0.6/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.0.6/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.0.5

Release Date: December 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.0.5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.0.5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.0.4

Release Date: August 2014 **Status:** Production

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.0.4/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.0.4/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.0.3

Release Date: July 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.0.3/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.0.3/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.0.2

Release Date: May 2014 **Status:** Production

Version and Download Information

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.0.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.0.2/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.0.1

Release Date: May 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.0.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.0.1/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5.0.0

Release Date: March 2014 **Status:** Production

Repository Type	Location	Repo file
Yum RHEL/CentOS/Oracle 6	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/5.0.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/6/x86_64/ cdh/cloudera-cdh5.repo</code>
Yum RHEL/CentOS/Oracle 5	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/5.0.0/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/redhat/5/x86_64/ cdh/cloudera-cdh5.repo</code>
Apt (Debian)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/debian/wheezy/ amd64/cdh/</code>	N/A
Zypper/YaST (SLES)	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/5/</code>	<code>https:// username:password@archive.cloudera.com/ p/cdh5/sles/11/x86_64/ cdh/cloudera-cdh5.repo</code>
Ubuntu	Precise: <code>https:// username:password@archive.cloudera.com/ p/cdh5/ubuntu/precise/ amd64/cdh</code>	N/A
Patched Tarball	<code>https:// username:password@archive.cloudera.com/ p/cdh5/cdh/5/</code>	N/A

CDH 5 Packaging and Tarball Information

Each CDH release series is made up of a collection of CDH project packages that are known to work together. The package version numbers of the CDH projects in each CDH release are listed in the following table.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.



Important: As of February 1, 2021, all downloads of CDH and Cloudera Manager require a username and password and use a modified URL. You must use the modified URL, including the username and password when downloading the repository contents described below. You may need to upgrade Cloudera Manager to a newer version that uses the modified URLs.

This can affect new installations, upgrades, adding new hosts to a cluster, adding a cluster and adding new services to a cluster.

For more information, see [Updating an existing CDH/Cloudera Manager deployment to access downloads with authentication](#) on page 99.

Examples of Versions

Cloudera packages are designed to be transparent and easy to customize. CDH 5 packages are labeled using the following format:

```
component-base_version+cdh_version+patch_level
```

where,

- `component-base_version` is the version of the open-source component included in the CDH package
- `cdh_version` is the version of the CDH package
- `patch_level` is the number of source commits applied on top of the base version forked from the Apache Hadoop branch. The list of source commits includes all backports and all non-functional changes such as CDH packaging and branding commits. Note that the number of commits does not indicate the number of functional changes or bug fixes in the release. For example, a commit may be used to amend a version number or make other non-functional changes. The list of actual patches that will match the count is found inside the CDH tarball's `../cloudera/patches/` directory. All of the source commits Cloudera has applied is also available in source form in that directory.

For example:

Package	Component	Branch	Base Version	CDH Version	Patch Level
hadoop-2.30+cdh5.10+795	hadoop	2.3	2.3.0	5.1.0	795
hue-3.6.0+cdh5.1.0+86	hue	3.6	3.6.0	5.1.0	86
parquet-1.25+cdh5.10+130	parquet	1.2	1.2.5	5.1.0	130

CDH Package Manifests

Both the CDH patched source and packages contain explicit information about Cloudera modifications. For example, in the patched source there is a top-level `cloudera` directory with:

- A `CHANGES.cloudera.txt` file that lists all the changes to the pristine source
- A `patches` directory that contains every patch Cloudera has applied to the pristine source. All Cloudera patches are released with an Apache 2.0 license.
- A `files` directory for files Cloudera created from scratch, such as man pages and configuration files. All Cloudera files are released with an Apache 2.0 license.
- A `README.cloudera` file that explains explicitly how to recreate the patches source from the pristine source.

Build and Release Numbering

If you are installing CDH 5 with a package manager, you will also see build and release information as part of the file name. The build and package release fields follow the patch level: for example,

```
hbase-0.98.1+cdh5.1.0+64-1.cdh5.1.0.p0.34.el6.x86_64.rpm.
```

The suffix `-1.cdh5.1.0.p0.34.el6.x86_64` represents:

- the base of the release field (1)
- the CDH release (`cdh5.1.0`)
- the customer patch identifier (`p0` — which will be 0 for all regular CDH releases but will increment for customer patches)
- the build number (34)
- the distribution (`el6` = RHEL/CentOs 6, `el5` = RHEL/Centos 5, `sles11` = SLES 11)
- the processor architecture (`x86_64`, `noarch`, `i386`, `amd64`). `noarch` means the packages are not architecture-specific.

External Documentation

Cloudera provides documentation for CDH as a whole, whether your CDH cluster is managed by Cloudera Manager or not. In addition, you may find it useful to refer to documentation for the individual components included in CDH. Where

Version and Download Information

possible, these links point to the main documentation for a project, in the Cloudera release archive. This ensures that you are looking at the correct documentation for the version of a project included in CDH. Otherwise, the links may point to the project's main site.

- [Apache Avro](#)
- [Apache Crunch](#)
- [Apache DataFu](#)
- [Apache Flume](#)
- [Apache Hadoop](#)
- [Apache HBase](#)
- [Apache Hive](#)
- [Hue](#)
- [Kite](#)
- [Apache Mahout](#)
- [Apache Oozie](#)
- [Apache Parquet](#)
- [Apache Pig](#)
- [Apache Sentry](#)
- [Apache Solr](#)
- [Apache Spark](#)
- [Apache Sqoop](#)
- [Apache Sqoop2](#)
- [Apache Whirr](#)
- [Apache ZooKeeper](#)

CDH 5.16.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.16.x release are listed in the following tables. For the same type of information for other CDH releases, see [CDH 5 Packaging and Tarball Information](#) on page 230. To view the overall release notes for CDH 5.x.x, see the [CDH 5 Release Notes](#) on page 1812.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.16.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.16.2+143	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.16.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.16.2.html	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.16.2.PDF.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.16.2+106	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.16.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.16.2.html	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.16.2.PDF.txt
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.16.2+29	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0+cdh5.16.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0+cdh5.16.2.html	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0+cdh5.16.2.PDF.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.7.0+cdh5.16.2+196	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.7.0+cdh5.16.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.7.0+cdh5.16.2.html	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.7.0+cdh5.16.2.PDF.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.16.2+2860	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.16.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.16.2.html	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.16.2.PDF.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-1.2.0+cdh5.16.2+489	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.2.0+cdh5.16.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.2.0+cdh5.16.2.html	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.2.0+cdh5.16.2.PDF.txt
		Tarball	Release notes	Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.16.2+77	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.5+cdh5.16.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.5+cdh5.16.2.html	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.5+cdh5.16.2.PDF.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.16.2+1450	https://us.apache.org/dist/cdh5/cdh/5/hive-1.1.0+cdh5.16.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hive-1.1.0+cdh5.16.2.html	https://us.apache.org/dist/cdh5/cdh/5/hive-1.1.0+cdh5.16.2.PDF.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-4.2.0+cdh5.16.2+8463	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hue-390-cdh5.16.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hue-390-cdh5.16.2.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hue-390-cdh5.16.2.PMS.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.12.0+cdh5.16.2+0	(none)	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/impala-200-cdh5.16.2.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/impala-200-cdh5.16.2.PMS.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.16.2+152	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kite-100-cdh5.16.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kite-100-cdh5.16.2.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kite-100-cdh5.16.2.PMS.txt
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.7.0+cdh5.16.2+0	(none)	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kudu-700-cdh5.16.2.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kudu-700-cdh5.16.2.PMS.txt
			Release notes	Changes
Llama	llama-1.0.0+cdh5.16.2+0	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama-100-cdh5.16.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama-100-cdh5.16.2.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama-100-cdh5.16.2.PMS.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.16.2+38	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout-09-cdh5.16.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout-09-cdh5.16.2.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout-09-cdh5.16.2.PMS.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.16.2+507	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-410-cdh5.16.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-410-cdh5.16.2.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-410-cdh5.16.2.PMS.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.16.2+200	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-150-cdh5.16.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-150-cdh5.16.2.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-150-cdh5.16.2.PMS.txt

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.16.2+22	https://us.apache.org/dist/parquet/cdh5/cdh/5/parquet-format-210cdh5162.tar.gz	https://us.apache.org/dist/parquet/cdh5/cdh/5/parquet-format-210cdh5162.html	https://us.apache.org/dist/parquet/cdh5/cdh/5/parquet-format-210cdh5162.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.16.2+119	https://us.apache.org/dist/pig/cdh5/cdh/5/pig-0120cdh5162.tar.gz	https://us.apache.org/dist/pig/cdh5/cdh/5/pig-0120cdh5162.html	https://us.apache.org/dist/pig/cdh5/cdh/5/pig-0120cdh5162.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.16.2+0	https://us.apache.org/dist/search/cdh5/cdh/5/search-100cdh5162.tar.gz	https://us.apache.org/dist/search/cdh5/cdh/5/search-100cdh5162.html	https://us.apache.org/dist/search/cdh5/cdh/5/search-100cdh5162.html#changes
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.16.2+573	https://us.apache.org/dist/sentry/cdh5/cdh/5/sentry-151cdh5162.tar.gz	https://us.apache.org/dist/sentry/cdh5/cdh/5/sentry-151cdh5162.html	https://us.apache.org/dist/sentry/cdh5/cdh/5/sentry-151cdh5162.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.16.2+535	https://us.apache.org/dist/solr/cdh5/cdh/5/solr-4103cdh5162.tar.gz	https://us.apache.org/dist/solr/cdh5/cdh/5/solr-4103cdh5162.html	https://us.apache.org/dist/solr/cdh5/cdh/5/solr-4103cdh5162.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.16.2+580	https://us.apache.org/dist/spark/cdh5/cdh/5/spark-160cdh5162.tar.gz	https://us.apache.org/dist/spark/cdh5/cdh/5/spark-160cdh5162.html	https://us.apache.org/dist/spark/cdh5/cdh/5/spark-160cdh5162.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.16.2+140	https://us.apache.org/dist/sqoop/cdh5/cdh/5/sqoop-146cdh5162.tar.gz	https://us.apache.org/dist/sqoop/cdh5/cdh/5/sqoop-146cdh5162.html	https://us.apache.org/dist/sqoop/cdh5/cdh/5/sqoop-146cdh5162.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.16.2+52	https://us.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop2-1995cdh5162.tar.gz	https://us.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop2-1995cdh5162.html	https://us.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop2-1995cdh5162.html#changes

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.16.2+27	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.16.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.16.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.16.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.16.2+159	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.16.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.16.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.16.2.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.16.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Avro	avro-1.7.6+cdh5.16.1+143	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.16.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.16.1.release.html	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.16.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.16.1+106	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.16.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.16.1.release.html	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.16.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.16.1+29	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.1.0+cdh5.16.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.1.0+cdh5.16.1.release.html	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.1.0+cdh5.16.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.7.0+cdh5.16.1+192	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.7.0+cdh5.16.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.7.0+cdh5.16.1.release.html	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.7.0+cdh5.16.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.16.1+2848	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.6.0+cdh5.16.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.6.0+cdh5.16.1.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.6.0+cdh5.16.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-1.2.0+cdh5.16.1			

Component	Package Version	Tarball	Release Notes	Changes File
	1+482	https://www.cloudera.com/resources/2016/05/16/120cdh5161.tar.gz	https://www.cloudera.com/resources/2016/05/16/120cdh5161.html	https://www.cloudera.com/resources/2016/05/16/120cdh5161.html#changes
		Tarball	Release notes	Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.16.1+77	https://www.cloudera.com/resources/2016/05/16/15cdh5161.tar.gz	https://www.cloudera.com/resources/2016/05/16/15cdh5161.html	https://www.cloudera.com/resources/2016/05/16/15cdh5161.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.16.1+1431	https://www.cloudera.com/resources/2016/05/16/hive110cdh5161.tar.gz	https://www.cloudera.com/resources/2016/05/16/hive110cdh5161.html	https://www.cloudera.com/resources/2016/05/16/hive110cdh5161.html#changes
		Tarball	Release notes	Changes
Hue	hue-4.2.0+cdh5.16.1+8441	https://www.cloudera.com/resources/2016/05/16/hue390cdh5161.tar.gz	https://www.cloudera.com/resources/2016/05/16/hue390cdh5161.html	https://www.cloudera.com/resources/2016/05/16/hue390cdh5161.html#changes
		Tarball	Release notes	Changes
Apache Impala	impala-2.12.0+cdh5.16.1+0	(none)	https://www.cloudera.com/resources/2016/05/16/imp200cdh5161.html	https://www.cloudera.com/resources/2016/05/16/imp200cdh5161.html#changes
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.16.1+151	https://www.cloudera.com/resources/2016/05/16/kite100cdh5161.tar.gz	https://www.cloudera.com/resources/2016/05/16/kite100cdh5161.html	https://www.cloudera.com/resources/2016/05/16/kite100cdh5161.html#changes
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.7.0+cdh5.16.1+0	(none)	https://www.cloudera.com/resources/2016/05/16/kud170cdh5161.html	https://www.cloudera.com/resources/2016/05/16/kud170cdh5161.html#changes
			Release notes	Changes
Llama	llama-1.0.0+cdh5.16.1+0	https://www.cloudera.com/resources/2016/05/16/llam100cdh5161.tar.gz	https://www.cloudera.com/resources/2016/05/16/llam100cdh5161.html	https://www.cloudera.com/resources/2016/05/16/llam100cdh5161.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Mahout	mahout-0.9+cdh5.16.1+38	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.16.1.tar.gz	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.16.1.html	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.16.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.16.1+503	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.16.1.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.16.1.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.16.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.16.1+200	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.16.1.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.16.1.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.16.1/CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.16.1+22	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.16.1.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.16.1.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.16.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.16.1+117	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.16.1.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.16.1.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.16.1/CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.16.1+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.16.1.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.16.1.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.16.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.16.1+559	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.16.1.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.16.1.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.16.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.16.1+532	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.16.1.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.16.1.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.16.1/CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Spark	spark-1.6.0+cdh5.16.1+577	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark160cdh5161.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark160cdh5161.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark160cdh5161Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.16.1+140	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop146cdh5161.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop146cdh5161.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop146cdh5161Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.16.1+52	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop21995cdh5161.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop21995cdh5161.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop21995cdh5161Changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.16.1+27	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr090cdh5161.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr090cdh5161.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr090cdh5161Changes.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.16.1+155	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper345cdh5161.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper345cdh5161.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper345cdh5161Changes.txt
		Tarball	Release notes	Changes

CDH 5.15.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.15.x release are listed in the following tables. For the same type of information for other CDH releases, see [CDH 5 Packaging and Tarball Information](#) on page 230. To view the overall release notes for CDH 5.x.x, see the [CDH 5 Release Notes](#) on page 1812.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.15.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.15.2+141	http://archive.cloudera.com/cdh5/cdh/5/avro176cdh5152.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/avro176cdh5152.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/avro176cdh5152.html#changes Changes
Apache Crunch	crunch-0.11.0+cdh5.15.2+105	http://archive.cloudera.com/cdh5/cdh/5/crunch0110cdh5152.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/crunch0110cdh5152.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/crunch0110cdh5152.html#changes Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.15.2+28	http://archive.cloudera.com/cdh5/cdh/5/datafu110cdh5152.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/datafu110cdh5152.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/datafu110cdh5152.html#changes Changes
Apache Flume	flume-ng-1.7.0+cdh5.15.2+190	http://archive.cloudera.com/cdh5/cdh/5/flume160cdh5152.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/flume160cdh5152.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/flume160cdh5152.html#changes Changes
Apache Hadoop	hadoop-2.6.0+cdh5.15.2+2843	http://archive.cloudera.com/cdh5/cdh/5/hadoop260cdh5152.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hadoop260cdh5152.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hadoop260cdh5152.html#changes Changes
Apache HBase	hbase-1.2.0+cdh5.15.2+479	http://archive.cloudera.com/cdh5/cdh/5/hbase120cdh5152.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hbase120cdh5152.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hbase120cdh5152.html#changes Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.15.2+75	http://archive.cloudera.com/cdh5/cdh/5/hbase15cdh5152.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hbase15cdh5152.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hbase15cdh5152.html#changes Changes
Apache Hive	hive-1.1.0+cdh5.15.2+1408	http://archive.cloudera.com/cdh5/cdh/5/hive110cdh5152.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hive110cdh5152.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hive110cdh5152.html#changes Changes
Hue	hue-4.2.0+cdh5.15.2+8443	http://archive.cloudera.com/cdh5/cdh/5/hue390cdh5152.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hue390cdh5152.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hue390cdh5152.html#changes Changes

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Impala	impala-2.12.0+cdh5.15.2+0	(none)	http://archive.cloudera.com/cdh5/cdh/5/impala200cdh520tar.gz	http://archive.cloudera.com/cdh5/cdh/5/impala200cdh520changes.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.15.2+148	http://archive.cloudera.com/cdh5/cdh/5/kite100cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/kite100cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/kite100cdh52changes.txt
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.7.0+cdh5.15.2+0	(none)	http://archive.cloudera.com/cdh5/cdh/5/kudu170cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/kudu170cdh52changes.txt
			Release notes	Changes
Llama	llama-1.0.0+cdh5.15.2+0	http://archive.cloudera.com/cdh5/cdh/5/llama100cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/llama100cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/llama100cdh52changes.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.15.2+37	http://archive.cloudera.com/cdh5/cdh/5/mahout09cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/mahout09cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/mahout09cdh52changes.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.15.2+500	http://archive.cloudera.com/cdh5/cdh/5/oozie410cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/oozie410cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/oozie410cdh52changes.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.15.2+198	http://archive.cloudera.com/cdh5/cdh/5/parquet150cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/parquet150cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/parquet150cdh52changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.15.2+21	http://archive.cloudera.com/cdh5/cdh/5/parquet-format210cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/parquet-format210cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/parquet-format210cdh52changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.15.2+115	http://archive.cloudera.com/cdh5/cdh/5/pig0120cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/pig0120cdh52tar.gz	http://archive.cloudera.com/cdh5/cdh/5/pig0120cdh52changes.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.15.2+0	http://archive.cloudera.com/cdh5/cdh/5/search-1.0.0+cdh5.15.2.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/search-1.0.0+cdh5.15.2.html	http://archive.cloudera.com/cdh5/cdh/5/search-1.0.0+cdh5.15.2.html
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.15.2+470	http://archive.cloudera.com/cdh5/cdh/5/sentry-1.5.1+cdh5.15.2.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/sentry-1.5.1+cdh5.15.2.html	http://archive.cloudera.com/cdh5/cdh/5/sentry-1.5.1+cdh5.15.2.html
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.15.2+530	http://archive.cloudera.com/cdh5/cdh/5/solr-4.10.3+cdh5.15.2.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/solr-4.10.3+cdh5.15.2.html	http://archive.cloudera.com/cdh5/cdh/5/solr-4.10.3+cdh5.15.2.html
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.15.2+574	http://archive.cloudera.com/cdh5/cdh/5/spark-1.6.0+cdh5.15.2.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/spark-1.6.0+cdh5.15.2.html	http://archive.cloudera.com/cdh5/cdh/5/spark-1.6.0+cdh5.15.2.html
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.15.2+137	http://archive.cloudera.com/cdh5/cdh/5/sqoop-1.4.6+cdh5.15.2.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/sqoop-1.4.6+cdh5.15.2.html	http://archive.cloudera.com/cdh5/cdh/5/sqoop-1.4.6+cdh5.15.2.html
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.15.2+51	http://archive.cloudera.com/cdh5/cdh/5/sqoop2-1.99.5+cdh5.15.2.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/sqoop2-1.99.5+cdh5.15.2.html	http://archive.cloudera.com/cdh5/cdh/5/sqoop2-1.99.5+cdh5.15.2.html
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.15.2+26	http://archive.cloudera.com/cdh5/cdh/5/whirr-0.9.0+cdh5.15.2.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/whirr-0.9.0+cdh5.15.2.html	http://archive.cloudera.com/cdh5/cdh/5/whirr-0.9.0+cdh5.15.2.html
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.15.2+153	http://archive.cloudera.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.15.2.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.15.2.html	http://archive.cloudera.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.15.2.html
		Tarball	Release notes	Changes

CDH 5.15.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.15.1+140	http://archive.cloudera.com/cdh5/cdh/5/avro1.7.6+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/avro1.7.6+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/avro1.7.6+cdh5.15.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.15.1+104	http://archive.cloudera.com/cdh5/cdh/5/crunch0.11.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/crunch0.11.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/crunch0.11.0+cdh5.15.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.15.1+27	http://archive.cloudera.com/cdh5/cdh/5/pig1.1.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/pig1.1.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/pig1.1.0+cdh5.15.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.7.0+cdh5.15.1+189	http://archive.cloudera.com/cdh5/cdh/5/flume1.6.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/flume1.6.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/flume1.6.0+cdh5.15.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.15.1+2822	http://archive.cloudera.com/cdh5/cdh/5/hadoop2.6.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/hadoop2.6.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/hadoop2.6.0+cdh5.15.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop MRv1	hadoop-0.20-mapreduce-2.6.0+cdh5.15.1+2822	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-1.2.0+cdh5.15.1+470	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.2.0+cdh5.15.1.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.2.0+cdh5.15.1-release.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.2.0+cdh5.15.1-changes.txt Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.15.1+74	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.5+cdh5.15.1.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.5+cdh5.15.1-release.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.5+cdh5.15.1-changes.txt Changes
Apache Hive	hive-1.1.0+cdh5.15.1+1395	http://archive.cloudera.com/cdh5/cdh/5/hive-1.1.0+cdh5.15.1.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hive-1.1.0+cdh5.15.1-release.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hive-1.1.0+cdh5.15.1-changes.txt Changes
Hue	hue-4.2.0+cdh5.15.1+8420	http://archive.cloudera.com/cdh5/cdh/5/hue-390+cdh5.15.1.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hue-390+cdh5.15.1-release.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hue-390+cdh5.15.1-changes.txt Changes
Apache Impala	impala-2.12.0+cdh5.15.1+0	(none)	http://archive.cloudera.com/cdh5/cdh/5/impala-2.12.0+cdh5.15.1-release.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/impala-2.12.0+cdh5.15.1-changes.txt Changes
Kite SDK	kite-1.0.0+cdh5.15.1+147	http://archive.cloudera.com/cdh5/cdh/5/kite-1.0.0+cdh5.15.1.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/kite-1.0.0+cdh5.15.1-release.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/kite-1.0.0+cdh5.15.1-changes.txt Changes
Apache Kudu	kudu-1.7.0+cdh5.15.1+0	(none)	http://archive.cloudera.com/cdh5/cdh/5/kudu-1.7.0+cdh5.15.1-release.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/kudu-1.7.0+cdh5.15.1-changes.txt Changes
Llama	llama-1.0.0+cdh5.15.1+0	http://archive.cloudera.com/cdh5/cdh/5/llama-1.0.0+cdh5.15.1.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/llama-1.0.0+cdh5.15.1-release.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/llama-1.0.0+cdh5.15.1-changes.txt Changes
Apache Mahout	mahout-0.9+cdh5.15.1+36	http://archive.cloudera.com/cdh5/cdh/5/mahout-0.9+cdh5.15.1.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/mahout-0.9+cdh5.15.1-release.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/mahout-0.9+cdh5.15.1-changes.txt Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.15.1+492	http://archive.cloudera.com/cdh5/cdh/5/oozie-4.1.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/oozie-4.1.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/oozie-4.1.0+cdh5.15.1.html
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.15.1+197	http://archive.cloudera.com/cdh5/cdh/5/parquet-1.5.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/parquet-1.5.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/parquet-1.5.0+cdh5.15.1.html
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.15.1+20	http://archive.cloudera.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.15.1.html
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.15.1+114	http://archive.cloudera.com/cdh5/cdh/5/pig-0.12.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/pig-0.12.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/pig-0.12.0+cdh5.15.1.html
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.15.1+0	http://archive.cloudera.com/cdh5/cdh/5/search-1.0.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/search-1.0.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/search-1.0.0+cdh5.15.1.html
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.15.1+458	http://archive.cloudera.com/cdh5/cdh/5/sentry-1.5.1+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/sentry-1.5.1+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/sentry-1.5.1+cdh5.15.1.html
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.15.1+529	http://archive.cloudera.com/cdh5/cdh/5/solr-4.10.3+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/solr-4.10.3+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/solr-4.10.3+cdh5.15.1.html
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.15.1+569	http://archive.cloudera.com/cdh5/cdh/5/spark-1.6.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/spark-1.6.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/spark-1.6.0+cdh5.15.1.html
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.15.1+136	http://archive.cloudera.com/cdh5/cdh/5/sqoop-1.4.6+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/sqoop-1.4.6+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/sqoop-1.4.6+cdh5.15.1.html
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop2	sqoop2-1.99.5+cdh5.15.1+49	http://archive.cloudera.com/cdh5/cdh/5/sqoop2-1.99.5+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/sqoop2-1.99.5+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/sqoop2-1.99.5+cdh5.15.1.html
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.15.1+25	http://archive.cloudera.com/cdh5/cdh/5/whirr-0.9.0+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/whirr-0.9.0+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/whirr-0.9.0+cdh5.15.1.html
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.15.1+149	http://archive.cloudera.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.15.1.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.15.1.html	http://archive.cloudera.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.15.1.html
		Tarball	Release notes	Changes

CDH 5.15.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.15.0+140	http://archive.cloudera.com/cdh5/cdh/5/avro-1.7.6+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/avro-1.7.6+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/avro-1.7.6+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.15.0+104	http://archive.cloudera.com/cdh5/cdh/5/crunch-0.11.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/crunch-0.11.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/crunch-0.11.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.15.0+27	http://archive.cloudera.com/cdh5/cdh/5/datafu-1.1.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/datafu-1.1.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/datafu-1.1.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.7.0+cdh5.15.0+185	http://archive.cloudera.com/cdh5/cdh/5/flume-ng-1.6.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/flume-ng-1.6.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/flume-ng-1.6.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.15.0+2783	http://archive.cloudera.com/cdh5/cdh/5/hadoop-2.6.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/hadoop-2.6.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/hadoop-2.6.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Hadoop MRv1	hadoop-0.20-mapreduce-2.6.	(none)	(none)	(none)

Component	Package Version	Tarball	Release Notes	Changes File
	0+cdh5.15.0+2783			

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-1.2.0+cdh5.15.0+461	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.2.0+cdh5.15.0.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.2.0+cdh5.15.0.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.2.0+cdh5.15.0.html#changes Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.15.0+74	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.5+cdh5.15.0.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.5+cdh5.15.0.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hbase-1.5+cdh5.15.0.html#changes Changes
Apache Hive	hive-1.1.0+cdh5.15.0+1373	http://archive.cloudera.com/cdh5/cdh/5/hive-1.1.0+cdh5.15.0.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hive-1.1.0+cdh5.15.0.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hive-1.1.0+cdh5.15.0.html#changes Changes
Hue	hue-4.2.0+cdh5.15.0+8368	http://archive.cloudera.com/cdh5/cdh/5/hue-390+cdh5.15.0.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/hue-390+cdh5.15.0.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/hue-390+cdh5.15.0.html#changes Changes
Apache Impala	impala-2.12.0+cdh5.15.0+0	(none)	http://archive.cloudera.com/cdh5/cdh/5/impala-2.12.0+cdh5.15.0.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/impala-2.12.0+cdh5.15.0.html#changes Changes
Kite SDK	kite-1.0.0+cdh5.15.0+147	http://archive.cloudera.com/cdh5/cdh/5/kite-1.0.0+cdh5.15.0.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/kite-1.0.0+cdh5.15.0.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/kite-1.0.0+cdh5.15.0.html#changes Changes
Apache Kudu	kudu-1.7.0+cdh5.15.0+0	(none)	http://archive.cloudera.com/cdh5/cdh/5/kudu-1.7.0+cdh5.15.0.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/kudu-1.7.0+cdh5.15.0.html#changes Changes
Llama	llama-1.0.0+cdh5.15.0+0	http://archive.cloudera.com/cdh5/cdh/5/llama-1.0.0+cdh5.15.0.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/llama-1.0.0+cdh5.15.0.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/llama-1.0.0+cdh5.15.0.html#changes Changes
Apache Mahout	mahout-0.9+cdh5.15.0+36	http://archive.cloudera.com/cdh5/cdh/5/mahout-0.9+cdh5.15.0.tar.gz Tarball	http://archive.cloudera.com/cdh5/cdh/5/mahout-0.9+cdh5.15.0.html Release notes	http://archive.cloudera.com/cdh5/cdh/5/mahout-0.9+cdh5.15.0.html#changes Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.15.0+487	http://archive.cloudera.com/cdh5/cdh/5/oozie-4.1.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/oozie-4.1.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/oozie-4.1.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.15.0+196	http://archive.cloudera.com/cdh5/cdh/5/parquet-1.5.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/parquet-1.5.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/parquet-1.5.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.15.0+20	http://archive.cloudera.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.15.0+114	http://archive.cloudera.com/cdh5/cdh/5/pig-0.12.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/pig-0.12.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/pig-0.12.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.15.0+0	http://archive.cloudera.com/cdh5/cdh/5/search-1.0.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/search-1.0.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/search-1.0.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.15.0+444	http://archive.cloudera.com/cdh5/cdh/5/sentry-1.5.1+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/sentry-1.5.1+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/sentry-1.5.1+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.15.0+524	http://archive.cloudera.com/cdh5/cdh/5/solr-4.10.3+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/solr-4.10.3+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/solr-4.10.3+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.15.0+547	http://archive.cloudera.com/cdh5/cdh/5/spark-1.6.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/spark-1.6.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/spark-1.6.0+cdh5.15.0.html
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.15.0+132	http://archive.cloudera.com/cdh5/cdh/5/sqoop-1.4.6+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/sqoop-1.4.6+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/sqoop-1.4.6+cdh5.15.0.html
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop2	sqoop2-1.99.5+cdh5.15.0+49	http://archive.cloudera.com/cdh5/cdh/5/sqoop2-1.99.5+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/sqoop2-1.99.5+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/sqoop2-1.99.5+cdh5.15.0.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.15.0+25	http://archive.cloudera.com/cdh5/cdh/5/whirr-0.9.0+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/whirr-0.9.0+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/whirr-0.9.0+cdh5.15.0.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.15.0+144	http://archive.cloudera.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.15.0.tar.gz	http://archive.cloudera.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.15.0.html	http://archive.cloudera.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.15.0.txt
		Tarball	Release notes	Changes

CDH 5.14.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.14.x release are listed in the following tables. For the same type of information for other CDH releases, see [CDH 5 Packaging and Tarball Information](#) on page 230. To view the overall release notes for CDH 5.x.x, see the [CDH 5 Release Notes](#) on page 1812.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.14.4 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.14.4+139	https://archive.apache.org/dist/avro/1.7.6+cdh5.14.4+139.tar.gz	https://archive.apache.org/dist/avro/1.7.6+cdh5.14.4+139/RELEASENOTES.html	https://archive.apache.org/dist/avro/1.7.6+cdh5.14.4+139/CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.14.4+106	https://archive.apache.org/dist/crunch/0.11.0+cdh5.14.4+106.tar.gz	https://archive.apache.org/dist/crunch/0.11.0+cdh5.14.4+106/RELEASENOTES.html	https://archive.apache.org/dist/crunch/0.11.0+cdh5.14.4+106/CHANGES.txt
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.14.4+29	https://archive.apache.org/dist/datafu/1.1.0+cdh5.14.4+29.tar.gz	https://archive.apache.org/dist/datafu/1.1.0+cdh5.14.4+29/RELEASENOTES.html	https://archive.apache.org/dist/datafu/1.1.0+cdh5.14.4+29/CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.7.0+cdh5.14.4+185	https://archive.apache.org/dist/flume-ng/1.7.0+cdh5.14.4+185.tar.gz	https://archive.apache.org/dist/flume-ng/1.7.0+cdh5.14.4+185/RELEASENOTES.html	https://archive.apache.org/dist/flume-ng/1.7.0+cdh5.14.4+185/CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.14.4+2785	https://archive.apache.org/dist/hadoop-core/2.6.0+cdh5.14.4+2785.tar.gz	https://archive.apache.org/dist/hadoop-core/2.6.0+cdh5.14.4+2785/RELEASENOTES.html	https://archive.apache.org/dist/hadoop-core/2.6.0+cdh5.14.4+2785/CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop MRv1	hadoop-0.20-mapreduce-2.6.0+cdh5.14.4+2785	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-1.2.0+cdh5.14.4+467	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.2.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.2.0+cdh5.14.4.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.2.0+cdh5.14.4.PDF.txt
		Tarball	Release notes	Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.14.4+75	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.5+cdh5.14.4.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.5+cdh5.14.4.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.5+cdh5.14.4.PDF.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.14.4+1360	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hive-1.1.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hive-1.1.0+cdh5.14.4.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hive-1.1.0+cdh5.14.4.PDF.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.14.4+7904	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hue-3.9.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hue-3.9.0+cdh5.14.4.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hue-3.9.0+cdh5.14.4.PDF.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.11.0+cdh5.14.4+0	(none)	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/impala-2.11.0+cdh5.14.4.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/impala-2.11.0+cdh5.14.4.PDF.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.14.4+149	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kite-1.0.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kite-1.0.0+cdh5.14.4.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kite-1.0.0+cdh5.14.4.PDF.txt
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.6.0+cdh5.14.4+0	(none)	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kudu-1.6.0+cdh5.14.4.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kudu-1.6.0+cdh5.14.4.PDF.txt
			Release notes	Changes
Llama	llama-1.0.0+cdh5.14.4+0	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/llama-1.0.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/llama-1.0.0+cdh5.14.4.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/llama-1.0.0+cdh5.14.4.PDF.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Mahout	mahout-0.9+cdh5.14.4+38	https://us.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.14.4.tar.gz	https://us.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.14.4.html	https://us.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.14.4.PM5.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.14.4+490	https://us.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.14.4.html	https://us.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.14.4.PM5.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.14.4+199	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.14.4.html	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.14.4.PM5.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.14.4+22	https://us.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.14.4.html	https://us.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.14.4.PM5.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.14.4+115	https://us.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.14.4.html	https://us.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.14.4.PM5.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.14.4+0	https://us.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.14.4.tar.gz	https://us.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.14.4.html	https://us.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.14.4.PM5.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.14.4+459	https://us.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.14.4.tar.gz	https://us.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.14.4.html	https://us.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.14.4.PM5.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.14.4+530	https://us.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.14.4.tar.gz	https://us.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.14.4.html	https://us.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.14.4.PM5.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Spark	spark-1.6.0+cdh5.14.4+565	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/spark160cdh5144tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/spark160cdh5144releasenotes.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/spark160cdh5144changes.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.14.4+134	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/sqoop146cdh5144tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/sqoop146cdh5144releasenotes.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/sqoop146cdh5144changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.14.4+49	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/sqoop21995cdh5144tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/sqoop21995cdh5144releasenotes.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/sqoop21995cdh5144changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.14.4+27	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/whirr090cdh5144tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/whirr090cdh5144releasenotes.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/whirr090cdh5144changes.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.14.4+144	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/zookeeper345cdh5144tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/zookeeper345cdh5144releasenotes.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/zookeeper345cdh5144changes.txt
		Tarball	Release notes	Changes

CDH 5.14.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.14.2+137	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/avro176cdh5142tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/avro176cdh5142releasenotes.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/avro176cdh5142changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.14.2+104	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/crunch0110cdh5142tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/crunch0110cdh5142releasenotes.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/crunch0110cdh5142changes.txt
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.14.2+26	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.14.2+26tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.14.2+26releasenotes.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.14.2+26changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		p/cdh5/cdh/5/cdh5-14.2+2748.tar.gz	p/cdh5/cdh/5/cdh5-14.2+2748.tar.gz	p/cdh5/cdh/5/cdh5-14.2+2748.tar.gz
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.7.0+cdh5.14.2+181	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume-ng-1.7.0+cdh5.14.2+181.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume-ng-1.7.0+cdh5.14.2+181.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume-ng-1.7.0+cdh5.14.2+181.tar.gz
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.14.2+2748	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.14.2+2748.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.14.2+2748.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.14.2+2748.tar.gz
		Tarball	Release notes	Changes
Apache Hadoop MRv1	hadoop-0.20-mapreduce-2.6.0+cdh5.14.2+2748	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-1.2.0+cdh5.14.2+456	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.14.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.14.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.14.2.PAGE.txt
		Tarball	Release notes	Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.14.2+73	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.14.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.14.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.14.2.PAGE.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.14.2+1344	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.14.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.14.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.14.2.PAGE.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.14.2+7875	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.14.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.14.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.14.2.PAGE.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.11.0+cdh5.14.2+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.11.0+cdh5.14.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.11.0+cdh5.14.2.PAGE.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.14.2+146	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.14.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.14.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.14.2.PAGE.txt
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.6.0+cdh5.14.2+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kudu-1.6.0+cdh5.14.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kudu-1.6.0+cdh5.14.2.PAGE.txt
			Release notes	Changes
Llama	llama-1.0.0+cdh5.14.2+0	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.14.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.14.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.14.2.PAGE.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Mahout	mahout-0.9+cdh5.14.2+35	https://www.apache.org/dist/mahout/mahout-0.9-cdh5.14.2.tar.gz	https://www.apache.org/dist/mahout/mahout-0.9-cdh5.14.2.html	https://www.apache.org/dist/mahout/mahout-0.9-cdh5.14.2.html#changes
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.14.2+484	https://www.apache.org/dist/oozie/oozie-4.1.0-cdh5.14.2.tar.gz	https://www.apache.org/dist/oozie/oozie-4.1.0-cdh5.14.2.html	https://www.apache.org/dist/oozie/oozie-4.1.0-cdh5.14.2.html#changes
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.14.2+196	https://www.apache.org/dist/parquet/parquet-1.5.0-cdh5.14.2.tar.gz	https://www.apache.org/dist/parquet/parquet-1.5.0-cdh5.14.2.html	https://www.apache.org/dist/parquet/parquet-1.5.0-cdh5.14.2.html#changes
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.14.2+19	https://www.apache.org/dist/parquet-format/parquet-format-2.1.0-cdh5.14.2.tar.gz	https://www.apache.org/dist/parquet-format/parquet-format-2.1.0-cdh5.14.2.html	https://www.apache.org/dist/parquet-format/parquet-format-2.1.0-cdh5.14.2.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.14.2+113	https://www.apache.org/dist/pig/pig-0.12.0-cdh5.14.2.tar.gz	https://www.apache.org/dist/pig/pig-0.12.0-cdh5.14.2.html	https://www.apache.org/dist/pig/pig-0.12.0-cdh5.14.2.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.14.2+0	https://www.apache.org/dist/search/search-1.0.0-cdh5.14.2.tar.gz	https://www.apache.org/dist/search/search-1.0.0-cdh5.14.2.html	https://www.apache.org/dist/search/search-1.0.0-cdh5.14.2.html#changes
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.14.2+444	https://www.apache.org/dist/sentry/sentry-1.5.1-cdh5.14.2.tar.gz	https://www.apache.org/dist/sentry/sentry-1.5.1-cdh5.14.2.html	https://www.apache.org/dist/sentry/sentry-1.5.1-cdh5.14.2.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.14.2+523	https://www.apache.org/dist/solr/solr-4.10.3-cdh5.14.2.tar.gz	https://www.apache.org/dist/solr/solr-4.10.3-cdh5.14.2.html	https://www.apache.org/dist/solr/solr-4.10.3-cdh5.14.2.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Spark	spark-1.6.0+cdh5.14.2+543	https://archive.apache.org/dist/parquet/cdh5/cdh/5/spark160cdh5142.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/spark160cdh5142.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/spark160cdh5142.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.14.2+130	https://archive.apache.org/dist/sqoop/cdh5/cdh/5/sqoop146cdh5142.tar.gz	https://archive.apache.org/dist/sqoop/cdh5/cdh/5/sqoop146cdh5142.html	https://archive.apache.org/dist/sqoop/cdh5/cdh/5/sqoop146cdh5142.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.14.2+47	https://archive.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop21995cdh5142.tar.gz	https://archive.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop21995cdh5142.html	https://archive.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop21995cdh5142.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.14.2+24	https://archive.apache.org/dist/whirr/cdh5/cdh/5/whirr090cdh5142.tar.gz	https://archive.apache.org/dist/whirr/cdh5/cdh/5/whirr090cdh5142.html	https://archive.apache.org/dist/whirr/cdh5/cdh/5/whirr090cdh5142.html#changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.14.2+142	https://archive.apache.org/dist/zookeeper/cdh5/cdh/5/zookeeper345cdh5142.tar.gz	https://archive.apache.org/dist/zookeeper/cdh5/cdh/5/zookeeper345cdh5142.html	https://archive.apache.org/dist/zookeeper/cdh5/cdh/5/zookeeper345cdh5142.html#changes
		Tarball	Release notes	Changes

CDH 5.14.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.14.0+137	https://archive.apache.org/dist/parquet/cdh5/cdh/5/avro176cdh5140.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/avro176cdh5140.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/avro176cdh5140.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.14.0+103	https://archive.apache.org/dist/parquet/cdh5/cdh/5/crunch0110cdh5140.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/crunch0110cdh5140.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/crunch0110cdh5140.html#changes
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.14.0+26	https://archive.apache.org/dist/parquet/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.14.0+26.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.14.0+26.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.14.0+26.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		<p>p/cdh5/cdh/5/cdh5.14.0.tar.gz</p> <p>Tarball</p>	<p>p/cdh5/cdh/5/cdh5.14.0.tar.gz</p> <p>Release notes</p>	<p>p/cdh5/cdh/5/cdh5.14.0.tar.gz</p> <p>Changes</p>
Apache Flume	flume-ng-1.7.0+cdh5.14.0+181	<p>https://archive.apache.org/dist/flume-ng/flume-ng-1.7.0+cdh5.14.0.tar.gz</p> <p>Tarball</p>	<p>https://archive.apache.org/dist/flume-ng/flume-ng-1.7.0+cdh5.14.0.tar.gz</p> <p>Release notes</p>	<p>https://archive.apache.org/dist/flume-ng/flume-ng-1.7.0+cdh5.14.0.tar.gz</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.14.0+2715	<p>https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.14.0.tar.gz</p> <p>Tarball</p>	<p>https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.14.0.tar.gz</p> <p>Release notes</p>	<p>https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.14.0.tar.gz</p> <p>Changes</p>
Apache Hadoop MRv1	hadoop-0.20-mapreduce-2.6.0+cdh5.14.0+2715	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-1.2.0+cdh5.14.0+440	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.14.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.14.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.14.0.PDF.txt
		Tarball	Release notes	Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.14.0+73	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.14.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.14.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.14.0.PDF.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.14.0+1331	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.14.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.14.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.14.0.PDF.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.14.0+7830	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.14.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.14.0.PDF.txt
			Release notes	Changes
Apache Impala	impala-2.11.0+cdh5.14.0+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.11.0+cdh5.14.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.11.0+cdh5.14.0.PDF.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.14.0+146	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.14.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.14.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.14.0.PDF.txt
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.6.0+cdh5.14.0+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kudu-1.6.0+cdh5.14.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kudu-1.6.0+cdh5.14.0.PDF.txt
			Release notes	Changes
Llama	llama-1.0.0+cdh5.14.0+0	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.14.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.14.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.14.0.PDF.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Mahout	mahout-0.9+cdh5.14.0+35	https://us.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9-cdh5.14.0.tar.gz	https://us.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9-cdh5.14.0.html	https://us.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9-cdh5.14.0.PM5.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.14.0+473	https://us.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.14.0.tar.gz	https://us.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.14.0.html	https://us.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.14.0.PM5.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.14.0+192	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0-cdh5.14.0.tar.gz	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0-cdh5.14.0.html	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0-cdh5.14.0.PM5.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.14.0+19	https://us.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.14.0.tar.gz	https://us.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.14.0.html	https://us.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.14.0.PM5.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.14.0+112	https://us.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0-cdh5.14.0.tar.gz	https://us.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0-cdh5.14.0.html	https://us.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0-cdh5.14.0.PM5.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.14.0+0	https://us.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0-cdh5.14.0.tar.gz	https://us.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0-cdh5.14.0.html	https://us.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0-cdh5.14.0.PM5.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.14.0+432	https://us.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1-cdh5.14.0.tar.gz	https://us.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1-cdh5.14.0.html	https://us.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1-cdh5.14.0.PM5.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.14.0+522	(none)	https://us.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3-cdh5.14.0.html	https://us.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3-cdh5.14.0.PM5.txt
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Spark	spark-1.6.0+cdh5.14.0+537	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.14.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.14.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.14.0.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.14.0+127	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.14.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.14.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.14.0.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.14.0+47	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.14.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.14.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.14.0.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.14.0+24	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.14.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.14.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.14.0.html#changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.14.0+136	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.14.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.14.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.14.0.html#changes
		Tarball	Release notes	Changes

CDH 5.13.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.13.x release are listed in the following tables. For the same type of information for other CDH releases, see [CDH 5 Packaging and Tarball Information](#) on page 230. To view the overall release notes for CDH 5.x.x, see the [CDH 5 Release Notes](#) on page 1812.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.13.3 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.13.3+139	https://archive.apache.org/dist/avro/1.7.6+cdh5.13.3+139.tar.gz	https://archive.apache.org/dist/avro/1.7.6+cdh5.13.3+139/RELEASENOTES.html	https://archive.apache.org/dist/avro/1.7.6+cdh5.13.3+139/CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.13.3+105	https://archive.apache.org/dist/crunch/0.11.0+cdh5.13.3+105.tar.gz	https://archive.apache.org/dist/crunch/0.11.0+cdh5.13.3+105/RELEASENOTES.html	https://archive.apache.org/dist/crunch/0.11.0+cdh5.13.3+105/CHANGES.txt
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.13.3+28	https://archive.apache.org/dist/datafu/1.1.0+cdh5.13.3+28.tar.gz	https://archive.apache.org/dist/datafu/1.1.0+cdh5.13.3+28/RELEASENOTES.html	https://archive.apache.org/dist/datafu/1.1.0+cdh5.13.3+28/CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.7.0+cdh5.13.3+180	https://archive.apache.org/dist/flume-ng/1.7.0+cdh5.13.3+180.tar.gz	https://archive.apache.org/dist/flume-ng/1.7.0+cdh5.13.3+180/RELEASENOTES.html	https://archive.apache.org/dist/flume-ng/1.7.0+cdh5.13.3+180/CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.13.3+2700	https://archive.apache.org/dist/hadoop-core/2.6.0+cdh5.13.3+2700.tar.gz	https://archive.apache.org/dist/hadoop-core/2.6.0+cdh5.13.3+2700/RELEASENOTES.html	https://archive.apache.org/dist/hadoop-core/2.6.0+cdh5.13.3+2700/CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop MRv1	hadoop-0.20-mapreduce-2.6.0+cdh5.13.3+2700	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-1.2.0+cdh5.13.3+441	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.13.3.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.13.3.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.13.3.PDF.txt
		Tarball	Release notes	Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.13.3+75	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.13.3.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.13.3.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.13.3.PDF.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.13.3+1307	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.13.3.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.13.3.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.13.3.PDF.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.13.3+7220	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.13.3.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.13.3.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.13.3.PDF.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.10.0+cdh5.13.3+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.10.0+cdh5.13.3.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.10.0+cdh5.13.3.PDF.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.13.3+149	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.13.3.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.13.3.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.13.3.PDF.txt
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.5.0+cdh5.13.3+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kudu-1.5.0+cdh5.13.3.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kudu-1.5.0+cdh5.13.3.PDF.txt
			Release notes	Changes
Llama	llama-1.0.0+cdh5.13.3+0	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.13.3.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.13.3.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.13.3.PDF.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Mahout	mahout-0.9+cdh5.13.3+38	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.13.3.tar.gz	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.13.3.html	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.13.3.PAOS.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.13.3+481	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.13.3.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.13.3.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.13.3.PAOS.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.13.3+198	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.13.3.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.13.3.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.13.3.PAOS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.13.3+22	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.13.3.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.13.3.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.13.3.PAOS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.13.3+115	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.13.3.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.13.3.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.13.3.PAOS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.13.3+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.13.3.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.13.3.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.13.3.PAOS.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.13.3+433	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.13.3.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.13.3.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.13.3.PAOS.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.13.3+525	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.13.3.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.13.3.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.13.3.PAOS.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Spark	spark-1.6.0+cdh5.13.3+543	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark160cdh5133.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark160cdh5133.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark160cdh5133CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.13.3+132	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop146cdh5133.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop146cdh5133.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop146cdh5133CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.13.3+50	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop21995cdh5133.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop21995cdh5133.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop21995cdh5133CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.13.3+27	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr090cdh5133.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr090cdh5133.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr090cdh5133CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.13.3+126	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper345cdh5133.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper345cdh5133.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper345cdh5133CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.13.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.13.2+138	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro176cdh5132.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro176cdh5132.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro176cdh5132CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.13.2+104	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch0110cdh5132.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch0110cdh5132.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch0110cdh5132CHANGES.txt
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.13.2+27	https://archive.apache.org/dist/parcels/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.13.2+27.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.13.2+27.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.13.2+27CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		<p>p/cdh5/cdh/5/cdh5-13.2+2687.tar.gz</p> <p>Tarball</p>	<p>p/cdh5/cdh/5/cdh5-13.2+2687.html</p> <p>Release notes</p>	<p>p/cdh5/cdh/5/cdh5-13.2+2687.txt</p> <p>Changes</p>
Apache Flume	flume-ng-1.7.0+cdh5.13.2+179	<p>https://archive.apache.org/dist/parcels/parcels/cdh5/cdh/5/flume-ng-1.7.0+cdh5.13.2+179.tar.gz</p> <p>Tarball</p>	<p>https://archive.apache.org/dist/parcels/parcels/cdh5/cdh/5/flume-ng-1.7.0+cdh5.13.2+179.html</p> <p>Release notes</p>	<p>https://archive.apache.org/dist/parcels/parcels/cdh5/cdh/5/flume-ng-1.7.0+cdh5.13.2+179.txt</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.13.2+2687	<p>https://archive.apache.org/dist/parcels/parcels/cdh5/cdh/5/hadoop-2.6.0+cdh5.13.2+2687.tar.gz</p> <p>Tarball</p>	<p>https://archive.apache.org/dist/parcels/parcels/cdh5/cdh/5/hadoop-2.6.0+cdh5.13.2+2687.html</p> <p>Release notes</p>	<p>https://archive.apache.org/dist/parcels/parcels/cdh5/cdh/5/hadoop-2.6.0+cdh5.13.2+2687.txt</p> <p>Changes</p>
Apache Hadoop MRv1	hadoop-0.20-mapreduce-2.6.0+cdh5.13.2+2687	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-1.2.0+cdh5.13.2+425	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.13.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.13.2-release.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.2.0+cdh5.13.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.13.2+74	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.13.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.13.2-release.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-1.5+cdh5.13.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.13.2+1297	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.13.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.13.2-release.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-1.1.0+cdh5.13.2-CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.13.2+7203	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.13.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.13.2-release.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.13.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.10.0+cdh5.13.2+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.10.0+cdh5.13.2-release.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.10.0+cdh5.13.2-CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.13.2+148	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.13.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.13.2-release.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.13.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.5.0+cdh5.13.2+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kudu-1.5.0+cdh5.13.2-release.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kudu-1.5.0+cdh5.13.2-CHANGES.txt
			Release notes	Changes
Llama	llama-1.0.0+cdh5.13.2+0	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.13.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.13.2-release.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.13.2-CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Mahout	mahout-0.9+cdh5.13.2+37	https://archive.apache.org/dist/mahout/mahout-0.9-cdh5.13.2.tar.gz	https://archive.apache.org/dist/mahout/mahout-0.9-cdh5.13.2.html	https://archive.apache.org/dist/mahout/mahout-0.9-cdh5.13.2.html#changes
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.13.2+477	https://archive.apache.org/dist/oozie/oozie-4.1.0-cdh5.13.2.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0-cdh5.13.2.html	https://archive.apache.org/dist/oozie/oozie-4.1.0-cdh5.13.2.html#changes
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.13.2+194	https://archive.apache.org/dist/parquet/parquet-1.5.0-cdh5.13.2.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0-cdh5.13.2.html	https://archive.apache.org/dist/parquet/parquet-1.5.0-cdh5.13.2.html#changes
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.13.2+21	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0-cdh5.13.2.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0-cdh5.13.2.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0-cdh5.13.2.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.13.2+114	https://archive.apache.org/dist/pig/pig-0.12.0-cdh5.13.2.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0-cdh5.13.2.html	https://archive.apache.org/dist/pig/pig-0.12.0-cdh5.13.2.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.13.2+0	https://archive.apache.org/dist/search/search-1.0.0-cdh5.13.2.tar.gz	https://archive.apache.org/dist/search/search-1.0.0-cdh5.13.2.html	https://archive.apache.org/dist/search/search-1.0.0-cdh5.13.2.html#changes
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.13.2+430	https://archive.apache.org/dist/sentry/sentry-1.5.1-cdh5.13.2.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1-cdh5.13.2.html	https://archive.apache.org/dist/sentry/sentry-1.5.1-cdh5.13.2.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.13.2+523	https://archive.apache.org/dist/solr/solr-4.10.3-cdh5.13.2.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3-cdh5.13.2.html	https://archive.apache.org/dist/solr/solr-4.10.3-cdh5.13.2.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Spark	spark-1.6.0+cdh5.13.2+539	https://archive.apache.org/dist/parquet/cdh5/cdh/5/spark160cdh5132.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/spark160cdh5132.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/spark160cdh5132CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.13.2+130	https://archive.apache.org/dist/sqoop/cdh5/cdh/5/sqoop146cdh5132.tar.gz	https://archive.apache.org/dist/sqoop/cdh5/cdh/5/sqoop146cdh5132.html	https://archive.apache.org/dist/sqoop/cdh5/cdh/5/sqoop146cdh5132CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.13.2+49	https://archive.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop21995cdh5132.tar.gz	https://archive.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop21995cdh5132.html	https://archive.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop21995cdh5132CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.13.2+26	https://archive.apache.org/dist/whirr/cdh5/cdh/5/whirr090cdh5132.tar.gz	https://archive.apache.org/dist/whirr/cdh5/cdh/5/whirr090cdh5132.html	https://archive.apache.org/dist/whirr/cdh5/cdh/5/whirr090cdh5132CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.13.2+125	https://archive.apache.org/dist/zookeeper/cdh5/cdh/5/zkeep345cdh5132.tar.gz	https://archive.apache.org/dist/zookeeper/cdh5/cdh/5/zkeep345cdh5132.html	https://archive.apache.org/dist/zookeeper/cdh5/cdh/5/zkeep345cdh5132CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.13.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.13.1+138	https://archive.apache.org/dist/parquet/cdh5/cdh/5/avro176cdh5131.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/avro176cdh5131.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/avro176cdh5131CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.13.1+104	https://archive.apache.org/dist/parquet/cdh5/cdh/5/crunch0110cdh5131.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/crunch0110cdh5131.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/crunch0110cdh5131CHANGES.txt
		Tarball	Release notes	Changes
Apache DataFu	pig-udf-datafu-1.1.0+cdh5.13.1+27	https://archive.apache.org/dist/parquet/cdh5/cdh/5/pigudfdatafu110cdh5131.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/pigudfdatafu110cdh5131.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/pigudfdatafu110cdh5131CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		p/cdh5/cdh/5/cdh5-13.1.1+2666.tar.gz	p/cdh5/cdh/5/cdh5-13.1.1+2666.tar.gz	p/cdh5/cdh/5/cdh5-13.1.1+2666.tar.gz
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.7.0+cdh5.13.1+174	https://archive.apache.org/dist/parcels/parcels/cdh5/cdh5-13.1.1+2666/flume-ng-1.7.0+cdh5.13.1+174.tar.gz	https://archive.apache.org/dist/parcels/parcels/cdh5/cdh5-13.1.1+2666/flume-ng-1.7.0+cdh5.13.1+174.tar.gz	https://archive.apache.org/dist/parcels/parcels/cdh5/cdh5-13.1.1+2666/flume-ng-1.7.0+cdh5.13.1+174.tar.gz
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.13.1+2666	https://archive.apache.org/dist/parcels/parcels/cdh5/cdh5-13.1.1+2666/hadoop-2.6.0+cdh5.13.1+2666.tar.gz	https://archive.apache.org/dist/parcels/parcels/cdh5/cdh5-13.1.1+2666/hadoop-2.6.0+cdh5.13.1+2666.tar.gz	https://archive.apache.org/dist/parcels/parcels/cdh5/cdh5-13.1.1+2666/hadoop-2.6.0+cdh5.13.1+2666.tar.gz
		Tarball	Release notes	Changes
Apache Hadoop MRv1	hadoop-0.20-mapreduce-2.6.0+cdh5.13.1+2666	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-1.2.0+cdh5.13.1+419	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.2.0+cdh5.13.1.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.2.0+cdh5.13.1.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.2.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase-Solr	hbase-solr-1.5+cdh5.13.1+74	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.5+cdh5.13.1.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.5+cdh5.13.1.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hbase-1.5+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.13.1+1283	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hive-1.1.0+cdh5.13.1.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hive-1.1.0+cdh5.13.1.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hive-1.1.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.13.1+7172	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hue-3.9.0+cdh5.13.1.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hue-3.9.0+cdh5.13.1.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/hue-3.9.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.10.0+cdh5.13.1+0	(none)	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/impala-2.10.0+cdh5.13.1.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/impala-2.10.0+cdh5.13.1/CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.13.1+148	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kite-1.0.0+cdh5.13.1.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kite-1.0.0+cdh5.13.1.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kite-1.0.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.5.0+cdh5.13.1+0	(none)	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kudu-1.5.0+cdh5.13.1.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/kudu-1.5.0+cdh5.13.1/CHANGES.txt
			Release notes	Changes
Llama	llama-1.0.0+cdh5.13.1+0	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/llama-1.0.0+cdh5.13.1.tar.gz	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/llama-1.0.0+cdh5.13.1.html	https://us.apache.org/dist/techtopcom/p/cdh5/cdh/5/llama-1.0.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Mahout	mahout-0.9+cdh5.13.1+37	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.13.1.tar.gz	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.13.1.html	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.13.1+464	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.13.1.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.13.1.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.13.1+194	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.13.1.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.13.1.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.13.1+21	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.13.1.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.13.1.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.13.1+113	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.13.1.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.13.1.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.13.1+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.13.1.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.13.1.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.13.1+417	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.13.1.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.13.1.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.13.1+522	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.13.1.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.13.1.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.13.1/CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Spark	spark-1.6.0+cdh5.13.1+536	https://archive.apache.org/dist/parquet/cdh5/cdh/5/spark160cdh5131.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/spark160cdh5131.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/spark160cdh5131CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.13.1+121	https://archive.apache.org/dist/sqoop/cdh5/cdh/5/sqoop146cdh5131.tar.gz	https://archive.apache.org/dist/sqoop/cdh5/cdh/5/sqoop146cdh5131.html	https://archive.apache.org/dist/sqoop/cdh5/cdh/5/sqoop146cdh5131CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.13.1+49	https://archive.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop21995cdh5131.tar.gz	https://archive.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop21995cdh5131.html	https://archive.apache.org/dist/sqoop2/cdh5/cdh/5/sqoop21995cdh5131CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.13.1+26	https://archive.apache.org/dist/whirr/cdh5/cdh/5/whirr090cdh5131.tar.gz	https://archive.apache.org/dist/whirr/cdh5/cdh/5/whirr090cdh5131.html	https://archive.apache.org/dist/whirr/cdh5/cdh/5/whirr090cdh5131CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.13.1+124	https://archive.apache.org/dist/zookeeper/cdh5/cdh/5/zk345cdh5131.tar.gz	https://archive.apache.org/dist/zookeeper/cdh5/cdh/5/zk345cdh5131.html	https://archive.apache.org/dist/zookeeper/cdh5/cdh/5/zk345cdh5131CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.13.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.13.0+135	https://archive.apache.org/dist/parquet/cdh5/cdh/5/avro176cdh5130.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/avro176cdh5130.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/avro176cdh5130CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.13.0+101	https://archive.apache.org/dist/parquet/cdh5/cdh/5/crunch0110cdh5130.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/crunch0110cdh5130.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/crunch0110cdh5130CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.13.0+24	https://archive.apache.org/dist/parquet/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.13.0+24.tar.gz	https://archive.apache.org/dist/parquet/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.13.0+24.html	https://archive.apache.org/dist/parquet/cdh5/cdh/5/pig-udf-datafu-1.1.0+cdh5.13.0+24CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		p/cdh5/cdh/5/cdh5-13.0+2639.tar.gz	p/cdh5/cdh/5/cdh5-13.0+2639.tar.gz	p/cdh5/cdh/5/cdh5-13.0+2639.tar.gz
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.7.0+cdh5.13.0+169	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume-ng-1.7.0+cdh5.13.0+169.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume-ng-1.7.0+cdh5.13.0+169.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume-ng-1.7.0+cdh5.13.0+169.tar.gz
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.13.0+2639	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.13.0+2639.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.13.0+2639.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.13.0+2639.tar.gz
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.13.0+2639	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.13.0+411	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.13.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.13.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.13.0.PMS.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.13.0+71	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.13.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.13.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.13.0.PMS.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.13.0+1269	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.13.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.13.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.13.0.PMS.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.13.0+7079	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.13.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.13.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.13.0.PMS.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.10.0+cdh5.13.0+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.10.0+cdh5.13.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.10.0+cdh5.13.0.PMS.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.13.0+145	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.13.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.13.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.13.0.PMS.txt
		Tarball	Release notes	Changes
Apache Kudu	kudu-1.5.0+cdh5.13.0+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/kudu-1.5.0+cdh5.13.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kudu-1.5.0+cdh5.13.0.PMS.txt
			Release notes	Changes
Llama	llama-1.0.0+cdh5.13.0+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.13.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.13.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.13.0.PMS.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Mahout	mahout-0.9+cdh5.13.0+34	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9-cdh5.13.0.tar.gz	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9-cdh5.13.0-release.html	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9-cdh5.13.0-CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.13.0+458	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.13.0.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.13.0-release.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.13.0-CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.13.0+191	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0-cdh5.13.0.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0-cdh5.13.0-release.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0-cdh5.13.0-CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.13.0+18	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.13.0.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.13.0-release.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.13.0-CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.13.0+110	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0-cdh5.13.0.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0-cdh5.13.0-release.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0-cdh5.13.0-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.13.0+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0-cdh5.13.0.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0-cdh5.13.0-release.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0-cdh5.13.0-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.13.0+410	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1-cdh5.13.0.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1-cdh5.13.0-release.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1-cdh5.13.0-CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.13.0+519	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3-cdh5.13.0.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3-cdh5.13.0-release.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3-cdh5.13.0-CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Spark	spark-1.6.0+cdh5.13.0+530	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.13.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.13.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.13.0.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.13.0+116	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.13.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.13.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.13.0.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.13.0+46	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.13.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.13.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.13.0.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.13.0+23	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.13.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.13.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.13.0.html#changes
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.13.0+118	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.13.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.13.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.13.0.html#changes
		Tarball	Release notes	Changes

CDH 5.12.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.12.x release are listed in the following tables. For the same type of information for other CDH releases, see [CDH 5 Packaging and Tarball Information](#) on page 230. To view the overall release notes for CDH 5.x.x, see the [CDH 5 Release Notes](#) on page 1812.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.12.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.12.2+134	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.12.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.12.2.html	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.12.2.PAGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.12.2+102	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.12.2.html	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.12.2.PAGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.12.2+25	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0+cdh5.12.2.html	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0+cdh5.12.2.PAGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.7.0+cdh5.12.2+173	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.7.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.7.0+cdh5.12.2.html	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.7.0+cdh5.12.2.PAGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.12.2+2573	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.12.2.html	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.12.2.PAGES.txt
		Tarball	Release notes	Changes
Hbase	hbase-1.2.0+cdh5.12.2+391	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.2.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.2.0+cdh5.12.2.html	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.2.0+cdh5.12.2.PAGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.12.2+72	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.5+cdh5.12.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.5+cdh5.12.2.html	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.5+cdh5.12.2.PAGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.12.2+1218	https://us.apache.org/dist/cdh5/cdh/5/hive-1.1.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hive-1.1.0+cdh5.12.2.html	https://us.apache.org/dist/cdh5/cdh/5/hive-1.1.0+cdh5.12.2.PAGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.12.2+6547	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.12.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.9.0+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Impala (Incubating)	impala-2.9.0+cdh5.12.2+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.9.0+cdh5.12.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-2.9.0+cdh5.12.2.html#changes
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.12.2+146	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.12.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-1.0.0+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.12.2+0	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.12.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.12.2+35	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/mahout-0.9+cdh5.12.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/mahout-0.9+cdh5.12.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/mahout-0.9+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.12.2+455	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/oozie-4.1.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/oozie-4.1.0+cdh5.12.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/oozie-4.1.0+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.12.2+188	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-1.5.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-1.5.0+cdh5.12.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-1.5.0+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.12.2+19	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.12.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.12.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.12.2.html#changes

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.12.2+111	https://archive.apache.org/dist/parcels/cdh5/cdh/5/pig-0.12.0+cdh5.12.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/pig-0.12.0+cdh5.12.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/pig-0.12.0+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.12.2+0	https://archive.apache.org/dist/parcels/cdh5/cdh/5/search-1.0.0+cdh5.12.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/search-1.0.0+cdh5.12.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/search-1.0.0+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.12.2+344	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.5.1+cdh5.12.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.5.1+cdh5.12.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.5.1+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.12.2+520	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.12.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.12.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.12.2+536	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.12.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.12.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.12.2+126	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.12.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.12.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.12.2+47	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.12.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.12.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.12.2.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.12.2+24	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.12.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.12.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.12.2.html#changes

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.12.2+122	https://archive.cloudera.com/p/cdh5/cdh/5/zoo-345-cdh5.12.2.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/zoo-345-cdh5.12.2.html	https://archive.cloudera.com/p/cdh5/cdh/5/zoo-345-cdh5.12.2.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.12.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.12.1+133	https://archive.cloudera.com/p/cdh5/cdh/5/av-176-cdh5.12.1.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/av-176-cdh5.12.1.html	https://archive.cloudera.com/p/cdh5/cdh/5/av-176-cdh5.12.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.12.1+101	https://archive.cloudera.com/p/cdh5/cdh/5/crunch-0110-cdh5.12.1.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/crunch-0110-cdh5.12.1.html	https://archive.cloudera.com/p/cdh5/cdh/5/crunch-0110-cdh5.12.1.CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.12.1+24	https://archive.cloudera.com/p/cdh5/cdh/5/datafu-110-cdh5.12.1.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/datafu-110-cdh5.12.1.html	https://archive.cloudera.com/p/cdh5/cdh/5/datafu-110-cdh5.12.1.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.7.0+cdh5.12.1+166	https://archive.cloudera.com/p/cdh5/cdh/5/flume-ng-160-cdh5.12.1.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/flume-ng-160-cdh5.12.1.html	https://archive.cloudera.com/p/cdh5/cdh/5/flume-ng-160-cdh5.12.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.12.1+2540	https://archive.cloudera.com/p/cdh5/cdh/5/hadoop-260-cdh5.12.1.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/hadoop-260-cdh5.12.1.html	https://archive.cloudera.com/p/cdh5/cdh/5/hadoop-260-cdh5.12.1.CHANGES.txt
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.12.1+2540	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.12.1+446	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.12.1.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.12.1.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.12.1+187	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.12.1.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.12.1.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.12.1+18	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.12.1.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.12.1.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.12.1+110	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.12.1.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.12.1.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.12.1+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.12.1.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.12.1.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.12.1+329	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.12.1.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.12.1.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.12.1+519	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.12.1.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.12.1.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.12.1+530	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.12.1.tar.gz	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.12.1.html	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.12.1+113	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.12.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.12.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.12.1+46	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.12.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.12.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.12.1+23	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.12.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.12.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.12.1+117	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.12.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.12.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.12.1Changes.txt
		Tarball	Release notes	Changes

CDH 5.12.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.12.0+132	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.12.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.12.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.12.0Changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.12.0+100	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.12.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.12.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.12.0Changes.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.12.0+23	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.12.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.12.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.12.0Changes.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.7.0+cdh5.12.0+163	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.7.0+cdh5.12.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.7.0+cdh5.12.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.7.0+cdh5.12.0Changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		<p>https://www.cloudera.com/p/cdh5/cdh/5/ha_260cdh5120.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/ha_260cdh5120.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/ha_260cdh5120.html#changes</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.12.0+2512	<p>https://www.cloudera.com/p/cdh5/cdh/5/ha_260cdh5120.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/ha_260cdh5120.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/ha_260cdh5120.html#changes</p> <p>Changes</p>
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.12.0+2512	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.12.0+442	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.12.0.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.12.0.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.12.0/PAGE.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.12.0+186	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.12.0.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.12.0.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.12.0/PAGE.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.12.0+17	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.12.0.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.12.0.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.12.0/PAGE.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.12.0+106	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.12.0.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.12.0.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.12.0/PAGE.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.12.0+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.12.0.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.12.0.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.12.0/PAGE.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.12.0+321	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.12.0.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.12.0.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.12.0/PAGE.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.12.0+513	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.12.0.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.12.0.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.12.0/PAGE.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.12.0+517	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.12.0.tar.gz	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.12.0.html	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.12.0/PAGE.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.12.0+111	https://archive.apache.org/dist/CDH/5/sqoop-1.4.6-cdh5.12.0.tar.gz	https://archive.apache.org/dist/CDH/5/release-notes/sqoop-1.4.6-cdh5.12.0.html	https://archive.apache.org/dist/CDH/5/changes/sqoop-1.4.6-cdh5.12.0.html
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.12.0+45	https://archive.apache.org/dist/CDH/5/sqoop2-1.99.5-cdh5.12.0.tar.gz	https://archive.apache.org/dist/CDH/5/release-notes/sqoop2-1.99.5-cdh5.12.0.html	https://archive.apache.org/dist/CDH/5/changes/sqoop2-1.99.5-cdh5.12.0.html
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.12.0+22	https://archive.apache.org/dist/CDH/5/whirr-0.9.0-cdh5.12.0.tar.gz	https://archive.apache.org/dist/CDH/5/release-notes/whirr-0.9.0-cdh5.12.0.html	https://archive.apache.org/dist/CDH/5/changes/whirr-0.9.0-cdh5.12.0.html
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.12.0+112	https://archive.apache.org/dist/CDH/5/zookeeper-3.4.5-cdh5.12.0.tar.gz	https://archive.apache.org/dist/CDH/5/release-notes/zookeeper-3.4.5-cdh5.12.0.html	https://archive.apache.org/dist/CDH/5/changes/zookeeper-3.4.5-cdh5.12.0.html
		Tarball	Release notes	Changes

CDH 5.11.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.11.0 release are listed in the following tables. For the same type of information for other CDH releases, see [CDH 5 Packaging and Tarball Information](#) on page 230. To view the overall release notes for CDH 5.x.x, see the [CDH 5 Release Notes](#) on page 1812.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.11.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.11.2+133	https://www.cloudera.com/p/cdh5/cdh/5/avro1.7.6cdh5.11.2.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/avro1.7.6cdh5.11.2.html	https://www.cloudera.com/p/cdh5/cdh/5/avro1.7.6cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.11.2+99	https://www.cloudera.com/p/cdh5/cdh/5/crunch0.11.0cdh5.11.2.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/crunch0.11.0cdh5.11.2.html	https://www.cloudera.com/p/cdh5/cdh/5/crunch0.11.0cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.11.2+24	https://www.cloudera.com/p/cdh5/cdh/5/datafu1.1.0cdh5.11.2.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/datafu1.1.0cdh5.11.2.html	https://www.cloudera.com/p/cdh5/cdh/5/datafu1.1.0cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.7.0+cdh5.11.2+165	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng1.7.0cdh5.11.2.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng1.7.0cdh5.11.2.html	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng1.7.0cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.11.2+2429	https://www.cloudera.com/p/cdh5/cdh/5/hadoop2.6.0cdh5.11.2.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop2.6.0cdh5.11.2.html	https://www.cloudera.com/p/cdh5/cdh/5/hadoop2.6.0cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.11.2+2429	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.11.2+441	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.11.2.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.11.2.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.11.2+187	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.11.2.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.11.2.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.11.2+18	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.11.2.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.11.2.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.11.2+108	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.11.2.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.11.2.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.11.2+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.11.2.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.11.2.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.11.2+318	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.11.2.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.11.2.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.11.2+514	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.11.2.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.11.2.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.11.2+519	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.11.2.tar.gz	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.11.2.html	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.11.2.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.11.2+112	https://archive.apache.org/dist/parcels/cdh5/cdh5/sqoop-1.4.6+cdh5.11.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh5/sqoop-1.4.6+cdh5.11.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh5/sqoop-1.4.6+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.11.2+46	https://archive.apache.org/dist/parcels/cdh5/cdh5/sqoop2-1.99.5+cdh5.11.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh5/sqoop2-1.99.5+cdh5.11.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh5/sqoop2-1.99.5+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.11.2+23	https://archive.apache.org/dist/parcels/cdh5/cdh5/whirr-0.9.0+cdh5.11.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh5/whirr-0.9.0+cdh5.11.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh5/whirr-0.9.0+cdh5.11.2.html#changes
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.11.2+117	https://archive.apache.org/dist/parcels/cdh5/cdh5/zoo-3.4.5+cdh5.11.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh5/zoo-3.4.5+cdh5.11.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh5/zoo-3.4.5+cdh5.11.2.html#changes
		Tarball	Release notes	Changes

CDH 5.11.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.11.1+132	https://archive.apache.org/dist/parcels/cdh5/cdh5/av-1.7.6+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh5/av-1.7.6+cdh5.11.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh5/av-1.7.6+cdh5.11.1.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.11.1+98	https://archive.apache.org/dist/parcels/cdh5/cdh5/crunch-0.11.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh5/crunch-0.11.0+cdh5.11.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh5/crunch-0.11.0+cdh5.11.1.html#changes
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.11.1+23	https://archive.apache.org/dist/parcels/cdh5/cdh5/datafu-1.1.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh5/datafu-1.1.0+cdh5.11.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh5/datafu-1.1.0+cdh5.11.1.html#changes
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.7.0+cdh5.11.1+160	https://archive.apache.org/dist/parcels/cdh5/cdh5/flume-ng-1.7.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh5/flume-ng-1.7.0+cdh5.11.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh5/flume-ng-1.7.0+cdh5.11.1.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.11.1.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.11.1-release-notes.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.11.1-changes.txt</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.11.1+2400	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.11.1.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.11.1-release-notes.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.11.1-changes.txt</p> <p>Changes</p>
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.11.1+2400	(none)	(none)	(none)

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.11.1+319	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.11.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.11.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.11.1-CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.11.1+70	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.11.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.11.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.11.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.11.1+1041	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.11.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.11.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.11.1-CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.11.1+5073	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.11.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.11.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.11.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.8.0+cdh5.11.1+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.8.0+cdh5.11.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.8.0+cdh5.11.1-CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.11.1+142	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.11.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.11.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.11.1-CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.11.1+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.11.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.11.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.11.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.11.1+33	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.11.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.11.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.11.1-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.11.1+431	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.11.1.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.11.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.11.1+185	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.11.1.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.11.1/CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.11.1+17	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.11.1.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.11.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.11.1+104	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.11.1.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.11.1/CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.11.1+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.11.1.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.11.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.11.1+310	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.11.1.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.11.1.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.11.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.11.1+502	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.11.1.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.11.1.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.11.1/CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.11.1+506	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.11.1.html	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.11.1/CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.11.1+111	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.11.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.11.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.11.1+44	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.11.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.11.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.11.1+22	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.11.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.11.1.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.11.1+111	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.11.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.11.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.11.1.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.11.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.11.0+131	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.11.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.11.0+97	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.11.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.11.0+22	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.11.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.7.0+cdh5.11.0+153	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.7.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.7.0+cdh5.11.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.7.0+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/impala-2.6.0+cdh5.11.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/impala-2.6.0+cdh5.11.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/impala-2.6.0+cdh5.11.0.PAHS.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.11.0+2345	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.11.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.11.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.11.0.PAHS.txt
		Tarball	Release notes	Changes
Hbase	hbase-1.2.0+cdh5.11.0+299	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.2.0+cdh5.11.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.2.0+cdh5.11.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.2.0+cdh5.11.0.PAHS.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.11.0+69	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.5+cdh5.11.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.5+cdh5.11.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.5+cdh5.11.0.PAHS.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.11.0+1013	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-1.1.0+cdh5.11.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-1.1.0+cdh5.11.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-1.1.0+cdh5.11.0.PAHS.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.11.0+5033	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.9.0+cdh5.11.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.9.0+cdh5.11.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.9.0+cdh5.11.0.PAHS.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.8.0+cdh5.11.0+0	(none)	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/impala-2.8.0+cdh5.11.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/impala-2.8.0+cdh5.11.0.PAHS.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.11.0+140	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/kite-1.0.0+cdh5.11.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/kite-1.0.0+cdh5.11.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/kite-1.0.0+cdh5.11.0.PAHS.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.11.0			

Component	Package Version	Tarball	Release Notes	Changes File
	0+0	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.0.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.0.0+cdh5.11.0.release.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.0.0+cdh5.11.0.changes.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.11.0+32	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.11.0.tar.gz	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.11.0.release.html	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.11.0.changes.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.11.0+426	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.11.0.release.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.11.0.changes.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.11.0+184	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.11.0.release.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.11.0.changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.11.0+16	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.11.0.release.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.11.0.changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.11.0+103	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.11.0.release.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.11.0.changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.11.0+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.11.0.release.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.11.0.changes.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.11.0+303	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.11.0.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.11.0.release.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.11.0.changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Solr	solr-4.10.3+cdh5.11.0+485	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.11.0.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.11.0+497	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.11.0.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.11.0+104	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.11.0.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.11.0+43	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.11.0.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.11.0+21	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.11.0.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.11.0+109	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.11.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.11.0.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.11.0.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.10.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.10.x release are listed in the following tables. For the same type of information for other CDH releases, see [CDH 5 Packaging and Tarball Information](#) on page 230. To view the overall release notes for CDH 5.x.x, see the [CDH 5 Release Notes](#) on page 1812.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.10.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.10.2+131	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6-cdh5.10.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6-cdh5.10.2.html	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6-cdh5.10.2.PKGS.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.10.2+98	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0-cdh5.10.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0-cdh5.10.2.html	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0-cdh5.10.2.PKGS.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.10.2+23	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0-cdh5.10.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0-cdh5.10.2.html	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0-cdh5.10.2.PKGS.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.10.2+103	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.6.0-cdh5.10.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.6.0-cdh5.10.2.html	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.6.0-cdh5.10.2.PKGS.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.10.2+2183	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0-cdh5.10.2.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0-cdh5.10.2.html	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0-cdh5.10.2.PKGS.txt
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.10.2+2183	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.10.2+303	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.10.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.10.2.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.10.2.PMS.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.10.2+70	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.10.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.10.2.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.10.2.PMS.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.10.2+945	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.10.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.10.2.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.10.2.PMS.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.10.2+4153	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.10.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.10.2.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.10.2.PMS.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.7.0+cdh5.10.2+0	(none)	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/impala-2.7.0+cdh5.10.2.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/impala-2.7.0+cdh5.10.2.PMS.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.10.2+141	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.10.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.10.2.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.10.2.PMS.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.10.2+0	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.10.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.10.2.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.10.2.PMS.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.10.2+33	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.10.2.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.10.2.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.10.2.PMS.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.10.2+411	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh5102.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh5102.release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh5102.PAGE.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.10.2+185	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh5102.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh5102.release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh5102.PAGE.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.10.2+17	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh5102.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh5102.release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh5102.PAGE.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.10.2+104	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh5102.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh5102.release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh5102.PAGE.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.10.2+0	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh5102.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh5102.release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh5102.PAGE.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.10.2+299	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh5102.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh5102.release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh5102.PAGE.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.10.2+503	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh5102.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh5102.release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh5102.PAGE.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.10.2+476	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh5102.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh5102.release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh5102.PAGE.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.10.2+111	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.10.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.10.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.10.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.10.2+45	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.10.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.10.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.10.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.10.2+22	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.10.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.10.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.10.2.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.10.2+108	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.10.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.10.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.10.2.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.10.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.10.1+130	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.10.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.10.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.10.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.10.1+97	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.10.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.10.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.10.1.CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.10.1+22	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.10.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.10.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.10.1.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.10.1+96	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.10.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.10.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.10.1.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		<p>p/cdh5/cdh/5/hadoop-2.6.0-cdh5.10.1.tar.gz</p> <p>Tarball</p>	<p>p/cdh5/cdh/5/hadoop-2.6.0-cdh5.10.1-release-notice.html</p> <p>Release notes</p>	<p>p/cdh5/cdh/5/hadoop-2.6.0-cdh5.10.1-changes.txt</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.10.1+2140	<p>https://www.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.10.1.tar.gz</p> <p>Tarball</p>	<p>https://www.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.10.1-release-notice.html</p> <p>Release notes</p>	<p>https://www.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.10.1-changes.txt</p> <p>Changes</p>
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.10.1+2140	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.10.1+276	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.10.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.10.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.10.1-CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.10.1+69	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.10.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.10.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.10.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.10.1+906	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.10.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.10.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.10.1-CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.10.1+4108	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.10.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.10.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.10.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.7.0+cdh5.10.1+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.7.0+cdh5.10.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.7.0+cdh5.10.1-CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.10.1+140	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.10.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.10.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.10.1-CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.10.1+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.10.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.10.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.10.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.10.1+32	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.10.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.10.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.10.1-CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.10.1+404	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh5101.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh5101.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh5101Changes.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.10.1+184	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh5101.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh5101.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh5101Changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.10.1+16	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh5101.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh5101.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh5101Changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.10.1+103	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh5101.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh5101.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh5101Changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.10.1+0	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh5101.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh5101.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh5101Changes.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.10.1+280	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh5101.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh5101.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh5101Changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.10.1+480	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh5101.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh5101.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh5101Changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.10.1+467	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh5101.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh5101.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh5101Changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.10.1+104	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.10.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.10.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.10.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.10.1+43	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.10.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.10.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.10.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.10.1+21	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.10.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.10.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.10.1.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.10.1+105	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.10.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.10.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.10.1.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.10.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.10.0+129	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.10.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.10.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.10.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.10.0+96	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.10.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.10.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.10.0.CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.10.0+21	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.10.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.10.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.10.0.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.10.0+86	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.10.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.10.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.10.0.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh5100.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh5100.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh5100.html#changes</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.10.0+2102	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh5100.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh5100.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh5100.html#changes</p> <p>Changes</p>
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.10.0+2102	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.10.0+236	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.10.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.10.0.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.2.0+cdh5.10.0.PDF.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.10.0+68	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.10.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.10.0.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.10.0.PDF.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.10.0+859	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.10.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.10.0.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.10.0.PDF.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.10.0+4082	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.10.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.10.0.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.10.0.PDF.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.7.0+cdh5.10.0+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.7.0+cdh5.10.0.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.7.0+cdh5.10.0.PDF.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.10.0+139	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.10.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.10.0.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.10.0.PDF.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.10.0+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.10.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.10.0.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.10.0.PDF.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.10.0+31	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.10.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.10.0.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.10.0.PDF.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.10.0+389	https://www.cloudera.com/p/cdh5/cdh/5/oozie-410cdh5100.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/oozie-410cdh5100.html	https://www.cloudera.com/p/cdh5/cdh/5/oozie-410cdh5100.PAGE.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.10.0+183	https://www.cloudera.com/p/cdh5/cdh/5/parquet-150cdh5100.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/parquet-150cdh5100.html	https://www.cloudera.com/p/cdh5/cdh/5/parquet-150cdh5100.PAGE.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.10.0+15	https://www.cloudera.com/p/cdh5/cdh/5/parquet-format-210cdh5100.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/parquet-format-210cdh5100.html	https://www.cloudera.com/p/cdh5/cdh/5/parquet-format-210cdh5100.PAGE.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.10.0+100	https://www.cloudera.com/p/cdh5/cdh/5/pig-0120cdh5100.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/pig-0120cdh5100.html	https://www.cloudera.com/p/cdh5/cdh/5/pig-0120cdh5100.PAGE.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.10.0+0	https://www.cloudera.com/p/cdh5/cdh/5/search-100cdh5100.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/search-100cdh5100.html	https://www.cloudera.com/p/cdh5/cdh/5/search-100cdh5100.PAGE.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.10.0+272	https://www.cloudera.com/p/cdh5/cdh/5/sentry-151cdh5100.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/sentry-151cdh5100.html	https://www.cloudera.com/p/cdh5/cdh/5/sentry-151cdh5100.PAGE.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.10.0+463	https://www.cloudera.com/p/cdh5/cdh/5/solr-4103cdh5100.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/solr-4103cdh5100.html	https://www.cloudera.com/p/cdh5/cdh/5/solr-4103cdh5100.PAGE.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.10.0+457	https://www.cloudera.com/p/cdh5/cdh/5/spark-160cdh5100.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/spark-160cdh5100.html	https://www.cloudera.com/p/cdh5/cdh/5/spark-160cdh5100.PAGE.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.10.0+89	https://archive.apache.org/dist/CDH/5/sqoop-1.4.6+cdh5.10.0.tar.gz	https://archive.apache.org/dist/CDH/5/release-notes/sqoop-1.4.6+cdh5.10.0.html	https://archive.apache.org/dist/CDH/5/changes/sqoop-1.4.6+cdh5.10.0.html
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.10.0+42	https://archive.apache.org/dist/CDH/5/sqoop2-1.99.5+cdh5.10.0.tar.gz	https://archive.apache.org/dist/CDH/5/release-notes/sqoop2-1.99.5+cdh5.10.0.html	https://archive.apache.org/dist/CDH/5/changes/sqoop2-1.99.5+cdh5.10.0.html
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.10.0+20	https://archive.apache.org/dist/CDH/5/whirr-0.9.0+cdh5.10.0.tar.gz	https://archive.apache.org/dist/CDH/5/release-notes/whirr-0.9.0+cdh5.10.0.html	https://archive.apache.org/dist/CDH/5/changes/whirr-0.9.0+cdh5.10.0.html
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.10.0+104	https://archive.apache.org/dist/CDH/5/zookeeper-3.4.5+cdh5.10.0.tar.gz	https://archive.apache.org/dist/CDH/5/release-notes/zookeeper-3.4.5+cdh5.10.0.html	https://archive.apache.org/dist/CDH/5/changes/zookeeper-3.4.5+cdh5.10.0.html
		Tarball	Release notes	Changes

CDH 5.9.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.9.x release are listed in the following tables. For the same type of information for other CDH releases, see [CDH 5 Packaging and Tarball Information](#) on page 230. To view the overall release notes for CDH 5.x.x, see the [CDH 5 Release Notes](#) on page 1812.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.9.3 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.9.3+119	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/avro-1.7.6+cdh5.9.3+119.tar.gz	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/avro-1.7.6+cdh5.9.3+119.html	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/avro-1.7.6+cdh5.9.3+119.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.9.3+97	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/crunch-0.11.0+cdh5.9.3+97.tar.gz	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/crunch-0.11.0+cdh5.9.3+97.html	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/crunch-0.11.0+cdh5.9.3+97.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.9.3+23	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/datafu-1.1.0+cdh5.9.3+23.tar.gz	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/datafu-1.1.0+cdh5.9.3+23.html	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/datafu-1.1.0+cdh5.9.3+23.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.9.3+103	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.3+103.tar.gz	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.3+103.html	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.3+103.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.9.3+1961	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.9.3+1961.tar.gz	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.9.3+1961.html	https://us.apache.org/dist/ahd/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.9.3+1961.txt
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.9.3+1961	(none)	(none)	(none)

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.9.3+377	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-410-cdh593.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-410-cdh593-release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-410-cdh593-CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.9.3+180	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-150-cdh593.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-150-cdh593-release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-150-cdh593-CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.9.3+17	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210-cdh593.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210-cdh593-release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210-cdh593-CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.9.3+104	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/pig-0120-cdh593.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/pig-0120-cdh593-release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/pig-0120-cdh593-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.9.3+0	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/search-100-cdh593.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/search-100-cdh593-release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/search-100-cdh593-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.9.3+291	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sentry-151-cdh593.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sentry-151-cdh593-release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sentry-151-cdh593-CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.9.3+500	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/solr-4103-cdh593.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/solr-4103-cdh593-release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/solr-4103-cdh593-CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.9.3+282	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/spark-160-cdh593.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/spark-160-cdh593-release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/spark-160-cdh593-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.9.3+112	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.9.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.9.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.9.3.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.9.3+45	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.9.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.9.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.9.3.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.9.3+22	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.9.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.9.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.9.3.html#changes
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.9.3+104	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.9.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.9.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.9.3.html#changes
		Tarball	Release notes	Changes

CDH 5.9.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.9.2+118	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.9.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.9.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.9.2.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.9.2+96	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.9.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.9.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.9.2.html#changes
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.9.2+22	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.9.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.9.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.9.2.html#changes
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.9.2+96	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.2.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh592.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh592.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh592CHANGES.txt</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.9.2+1942	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh592.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh592.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh592CHANGES.txt</p> <p>Changes</p>
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.9.2+1942	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.9.2+271	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.9.2.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.9.2.release.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.9.2+69	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.9.2.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.9.2.release.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.9.2+855	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.9.2.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.9.2.release.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.9.2+3324	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.9.2.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.9.2.release.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.7.0+cdh5.9.2+0	(none)	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/impala-2.7.0+cdh5.9.2.release.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/impala-2.7.0+cdh5.9.2.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.9.2+140	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.9.2.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.9.2.release.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.9.2+0	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.9.2.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.9.2.release.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.9.2+32	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.9.2.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.9.2.release.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.9.2+369	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh592.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh592.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh592.PMS.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.9.2+178	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh592.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh592.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh592.PMS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.9.2+16	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh592.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh592.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh592.PMS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.9.2+102	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh592.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh592.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh592.PMS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.9.2+0	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh592.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh592.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh592.PMS.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.9.2+277	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh592.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh592.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh592.PMS.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.9.2+484	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh592.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh592.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh592.PMS.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.9.2+264	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh592.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh592.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh592.PMS.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.9.2+109	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.9.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.9.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.9.2+42	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.9.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.9.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.9.2+21	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.9.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.9.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.9.2+100	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.9.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.9.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.9.2.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.9.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.9.1+117	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.9.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.9.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.9.1+95	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.9.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.9.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.9.1+21	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.9.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.9.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.9.1+80	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.9.1+1889.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.9.1+1889.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.9.1+1889.tar.gz
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.9.1+1889	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.9.1+1889.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.9.1+1889.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.9.1+1889.tar.gz
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.9.1+1889	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.9.1+222	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.9.1.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.9.1.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.9.1+68	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.9.1.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.9.1.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.9.1+793	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.9.1.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.9.1.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.9.1+3282	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.9.1.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.9.1.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.7.0+cdh5.9.1+0	(none)	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/impala-2.7.0+cdh5.9.1.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/impala-2.7.0+cdh5.9.1.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.9.1+139	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.9.1.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.9.1.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.9.1+0	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.9.1.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.9.1.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.9.1+31	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.9.1.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.9.1.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.9.1+353	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/oozie-410-cdh591.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/oozie-410-cdh591-release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/oozie-410-cdh591-CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.9.1+177	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-150-cdh591.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-150-cdh591-release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-150-cdh591-CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.9.1+15	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210-cdh591.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210-cdh591-release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210-cdh591-CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.9.1+98	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/pig-0120-cdh591.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/pig-0120-cdh591-release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/pig-0120-cdh591-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.9.1+0	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/search-100-cdh591.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/search-100-cdh591-release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/search-100-cdh591-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.9.1+266	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/sentry-151-cdh591.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/sentry-151-cdh591-release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/sentry-151-cdh591-CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.9.1+460	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/solr-4103-cdh591.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/solr-4103-cdh591-release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/solr-4103-cdh591-CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.9.1+245	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/spark-160-cdh591.tar.gz	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/spark-160-cdh591-release.html	https://usapocdn.blob.core.windows.net/p/cdh5/cdh/5/spark-160-cdh591-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.9.1+85	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.9.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.9.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.9.1+41	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.9.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.9.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.9.1+20	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.9.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.9.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.9.1+99	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.9.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.9.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.9.1.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.9.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.9.0+115	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.9.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.9.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.9.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.9.0+93	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.9.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.9.0.CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.9.0+20	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.9.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.9.0.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.9.0+74	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.9.0.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		https://www.cloudera.com/p/cdh5/cdh/5/160cdh590.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/160cdh590.html	https://www.cloudera.com/p/cdh5/cdh/5/160cdh590CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.9.0+1799	https://www.cloudera.com/p/cdh5/cdh/5/160cdh590.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/160cdh590.html	https://www.cloudera.com/p/cdh5/cdh/5/160cdh590CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.9.0+1799	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
HBase	hbase-1.2.0+cdh5.9.0+205	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.2.0+cdh5.9.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.2.0+cdh5.9.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.2.0+cdh5.9.0.PMS.txt
		Tarball	Release notes	Changes
HBase-solr	hbase-solr-1.5+cdh5.9.0+67	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.5+cdh5.9.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.5+cdh5.9.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.5+cdh5.9.0.PMS.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.9.0+752	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-1.1.0+cdh5.9.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-1.1.0+cdh5.9.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-1.1.0+cdh5.9.0.PMS.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.9.0+3235	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.9.0+cdh5.9.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.9.0+cdh5.9.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.9.0+cdh5.9.0.PMS.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.7.0+cdh5.9.0+0	(none)	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/impala-2.7.0+cdh5.9.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/impala-2.7.0+cdh5.9.0.PMS.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.9.0+138	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/kite-1.0.0+cdh5.9.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/kite-1.0.0+cdh5.9.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/kite-1.0.0+cdh5.9.0.PMS.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.9.0+0	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/llama-1.0.0+cdh5.9.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/llama-1.0.0+cdh5.9.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/llama-1.0.0+cdh5.9.0.PMS.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.9.0+30	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/mahout-0.9+cdh5.9.0.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/mahout-0.9+cdh5.9.0.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/mahout-0.9+cdh5.9.0.PMS.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.9.0+349	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.9.0.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.9.0.html#changes
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.9.0+176	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.9.0.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.9.0.html#changes
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.9.0+14	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.9.0.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.9.0.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.9.0+95	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.9.0.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.9.0.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.9.0+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.9.0.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.9.0.html#changes
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.9.0+261	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.9.0.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.9.0.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.9.0.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.9.0+458	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.9.0.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.9.0.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.9.0.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.9.0+229	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.9.0.html	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.9.0.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.9.0+77	https://archive.apache.org/dist/CDH/5/sqoop-1.4.6+cdh5.9.0.tar.gz	https://archive.apache.org/dist/CDH/5/sqoop-1.4.6+cdh5.9.0.html	https://archive.apache.org/dist/CDH/5/sqoop-1.4.6+cdh5.9.0.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.9.0+40	https://archive.apache.org/dist/CDH/5/sqoop2-1.99.5+cdh5.9.0.tar.gz	https://archive.apache.org/dist/CDH/5/sqoop2-1.99.5+cdh5.9.0.html	https://archive.apache.org/dist/CDH/5/sqoop2-1.99.5+cdh5.9.0.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.9.0+19	https://archive.apache.org/dist/CDH/5/whirr-0.9.0+cdh5.9.0.tar.gz	https://archive.apache.org/dist/CDH/5/whirr-0.9.0+cdh5.9.0.html	https://archive.apache.org/dist/CDH/5/whirr-0.9.0+cdh5.9.0.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.9.0+98	https://archive.apache.org/dist/CDH/5/zoo-3.4.5+cdh5.9.0.tar.gz	https://archive.apache.org/dist/CDH/5/zoo-3.4.5+cdh5.9.0.html	https://archive.apache.org/dist/CDH/5/zoo-3.4.5+cdh5.9.0.txt
		Tarball	Release notes	Changes

CDH 5.8.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.8.x release are listed in the following table. For the same type of information for other CDH releases, see [CDH 5 Packaging and Tarball Information](#) on page 230. To view the overall release notes for CDH 5.x.x, see the [CDH 5 Release Notes](#) on page 1812.



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.8.5 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.8.5+118	https://archive.apache.org/dist/avro/1.7.6+cdh5.8.5+118/avro-1.7.6+cdh5.8.5+118.tar.gz	https://archive.apache.org/dist/avro/1.7.6+cdh5.8.5+118/avro-1.7.6+cdh5.8.5+118.html	https://archive.apache.org/dist/avro/1.7.6+cdh5.8.5+118/avro-1.7.6+cdh5.8.5+118.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.8.5+97	https://archive.apache.org/dist/crunch/0.11.0+cdh5.8.5+97/crunch-0.11.0+cdh5.8.5+97.tar.gz	https://archive.apache.org/dist/crunch/0.11.0+cdh5.8.5+97/crunch-0.11.0+cdh5.8.5+97.html	https://archive.apache.org/dist/crunch/0.11.0+cdh5.8.5+97/crunch-0.11.0+cdh5.8.5+97.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.8.5+23	https://archive.apache.org/dist/datafu/1.1.0+cdh5.8.5+23/datafu-1.1.0+cdh5.8.5+23.tar.gz	https://archive.apache.org/dist/datafu/1.1.0+cdh5.8.5+23/datafu-1.1.0+cdh5.8.5+23.html	https://archive.apache.org/dist/datafu/1.1.0+cdh5.8.5+23/datafu-1.1.0+cdh5.8.5+23.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.8.5+94	https://archive.apache.org/dist/flume-ng/1.6.0+cdh5.8.5+94/flume-ng-1.6.0+cdh5.8.5+94.tar.gz	https://archive.apache.org/dist/flume-ng/1.6.0+cdh5.8.5+94/flume-ng-1.6.0+cdh5.8.5+94.html	https://archive.apache.org/dist/flume-ng/1.6.0+cdh5.8.5+94/flume-ng-1.6.0+cdh5.8.5+94.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.8.5+1850	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.8.5+1850.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.8.5+1850.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.8.5+1850.txt
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.8.5+1850	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.8.5+263	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.8.5.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.8.5.rn.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.8.5.changes.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.8.5+70	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.8.5.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.8.5.rn.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.8.5.changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.8.5+774	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.8.5.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.8.5.rn.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.8.5.changes.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.8.5+2690	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.8.5.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.8.5.rn.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.8.5.changes.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.6.0+cdh5.8.5+0	(none)	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/impala-2.6.0+cdh5.8.5.rn.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/impala-2.6.0+cdh5.8.5.changes.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.8.5+141	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.8.5.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.8.5.rn.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.8.5.changes.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.8.5+0	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.8.5.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.8.5.rn.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.8.5.changes.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.8.5+32	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.8.5.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.8.5.rn.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-123456789012-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.8.5.changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.8.5+327	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.8.5.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.8.5.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.8.5.PAUS.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.8.5+179	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.8.5.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.8.5.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.8.5.PAUS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.8.5+17	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.8.5.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.8.5.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.8.5.PAUS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.8.5+95	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.8.5.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.8.5.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.8.5.PAUS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.8.5+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.8.5.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.8.5.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.8.5.PAUS.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.8.5+267	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.8.5.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.8.5.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.8.5.PAUS.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.8.5+464	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.8.5.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.8.5.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.8.5.PAUS.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.8.5+256	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.6.0+cdh5.8.5.tar.gz	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.6.0+cdh5.8.5.html	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.6.0+cdh5.8.5.PAUS.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.8.5+107	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop146cdh585.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop146cdh585.html	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop146cdh585.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.8.5+43	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop21995cdh585.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop21995cdh585.html	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop21995cdh585.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.8.5+22	https://archive.cloudera.com/p/cdh5/cdh/5/whirr090cdh585.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/whirr090cdh585.html	https://archive.cloudera.com/p/cdh5/cdh/5/whirr090cdh585.html#changes
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.8.5+98	https://archive.cloudera.com/p/cdh5/cdh/5/zoo345cdh585.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/zoo345cdh585.html	https://archive.cloudera.com/p/cdh5/cdh/5/zoo345cdh585.html#changes
		Tarball	Release notes	Changes

CDH 5.8.4 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.8.4+117	https://archive.cloudera.com/p/cdh5/cdh/5/avro176cdh584.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/avro176cdh584.html	https://archive.cloudera.com/p/cdh5/cdh/5/avro176cdh584.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.8.4+96	https://archive.cloudera.com/p/cdh5/cdh/5/crunch0110cdh584.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/crunch0110cdh584.html	https://archive.cloudera.com/p/cdh5/cdh/5/crunch0110cdh584.html#changes
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.8.4+22	https://archive.cloudera.com/p/cdh5/cdh/5/datafu110cdh584.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/datafu110cdh584.html	https://archive.cloudera.com/p/cdh5/cdh/5/datafu110cdh584.html#changes
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.8.4+89	https://archive.cloudera.com/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.8.4+89.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.8.4+89.html	https://archive.cloudera.com/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.8.4+89.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/impala-1.6.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/impala-1.6.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/impala-1.6.0+cdh5.8.4.tar.gz
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.8.4+1801	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.8.4.tar.gz
		Tarball	Release notes	Changes
Hbase	hbase-1.2.0+cdh5.8.4+219	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.8.4.tar.gz
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.8.4+69	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.8.4.tar.gz
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.8.4+723	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.8.4.tar.gz
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.8.4+2609	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.8.4.tar.gz
		Tarball	Release notes	Changes
Apache Impala	impala-2.6.0+cdh5.8.4+0	(none)	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/impala-2.6.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/impala-2.6.0+cdh5.8.4.tar.gz
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.8.4+140	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.8.4.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-101123180201-us-east-1-bucket-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.8.4.tar.gz
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.8.4			

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
	4+0	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/lan-1.0.0-cdh5.8.4.tar.gz	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/lan-1.0.0-cdh5.8.4.html	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/lan-1.0.0-cdh5.8.4.PM5.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.8.4+31	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/mahout-0.9-cdh5.8.4.tar.gz	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/mahout-0.9-cdh5.8.4.html	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/mahout-0.9-cdh5.8.4.PM5.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.8.4+315	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/oozie-4.1.0-cdh5.8.4.tar.gz	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/oozie-4.1.0-cdh5.8.4.html	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/oozie-4.1.0-cdh5.8.4.PM5.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.8.4+178	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/parquet-1.5.0-cdh5.8.4.tar.gz	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/parquet-1.5.0-cdh5.8.4.html	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/parquet-1.5.0-cdh5.8.4.PM5.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.8.4+16	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.8.4.tar.gz	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.8.4.html	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.8.4.PM5.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.8.4+90	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/pig-0.12.0-cdh5.8.4.tar.gz	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/pig-0.12.0-cdh5.8.4.html	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/pig-0.12.0-cdh5.8.4.PM5.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.8.4+0	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/search-1.0.0-cdh5.8.4.tar.gz	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/search-1.0.0-cdh5.8.4.html	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/search-1.0.0-cdh5.8.4.PM5.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.8.4+256	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/sentry-1.5.1-cdh5.8.4.tar.gz	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/sentry-1.5.1-cdh5.8.4.html	https://us.apache.org/dist/mirrors/p/cdh5/cdh/5/sentry-1.5.1-cdh5.8.4.PM5.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Solr	solr-4.10.3+cdh5.8.4+441	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.8.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.8.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.8.4.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.8.4+243	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.8.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.8.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.8.4.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.8.4+100	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.8.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.8.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.8.4.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.8.4+42	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.8.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.8.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.8.4.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.8.4+21	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.8.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.8.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.8.4.html#changes
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.8.4+97	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.8.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.8.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.8.4.html#changes
		Tarball	Release notes	Changes

CDH 5.8.3 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.8.3+115	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.8.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.8.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.8.3.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.8.3+94	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.8.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.8.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.8.3.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		<p>https://www.cloudera.com/p/cdh5/cdh/5/cdh5.0.110-cdh5.8.3.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/cdh5.0.110-cdh5.8.3-release-notes.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/cdh5.0.110-cdh5.8.3-changes.txt</p> <p>Changes</p>
Datafu	pig-udf-datafu-1.1.0+cdh5.8.3+21	<p>https://www.cloudera.com/p/cdh5/cdh/5/datafu-1.10-cdh5.8.3.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/datafu-1.10-cdh5.8.3-release-notes.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/datafu-1.10-cdh5.8.3-changes.txt</p> <p>Changes</p>
Flume-ng	flume-ng-1.6.0+cdh5.8.3+75	<p>https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.6.0-cdh5.8.3.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.6.0-cdh5.8.3-release-notes.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.6.0-cdh5.8.3-changes.txt</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.8.3+1718	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.8.3.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.8.3-release-notes.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.8.3-changes.txt</p> <p>Changes</p>
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.8.3+1718	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.8.3+314	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.8.3+314.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.8.3+314.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.8.3+314.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.8.3+177	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.8.3+177.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.8.3+177.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.8.3+177.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.8.3+15	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.8.3+15.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.8.3+15.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.8.3+15.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.8.3+87	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.8.3+87.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.8.3+87.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.8.3+87.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.8.3+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.8.3+0.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.8.3+0.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.8.3+0.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.8.3+251	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.8.3+251.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.8.3+251.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.8.3+251.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.8.3+439	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.8.3+439.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.8.3+439.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.8.3+439.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.8.3+232	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.8.3+232.tar.gz	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.8.3+232.html	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.8.3+232.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.8.3+76	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6-cdh5.8.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6-cdh5.8.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6-cdh5.8.3CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.8.3+41	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5-cdh5.8.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5-cdh5.8.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5-cdh5.8.3CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.8.3+20	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0-cdh5.8.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0-cdh5.8.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0-cdh5.8.3CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.8.3+96	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5-cdh5.8.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5-cdh5.8.3.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5-cdh5.8.3CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.8.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.8.2+114	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6-cdh5.8.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6-cdh5.8.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6-cdh5.8.2CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.8.2+93	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0-cdh5.8.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0-cdh5.8.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0-cdh5.8.2CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.8.2+20	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0-cdh5.8.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0-cdh5.8.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0-cdh5.8.2CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.8.2+72	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0-cdh5.8.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0-cdh5.8.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0-cdh5.8.2CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.8.2+1693.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.8.2+1693.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.8.2+1693.html</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.8.2+1693	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.8.2+1693.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.8.2+1693.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.8.2+1693.html</p> <p>Changes</p>
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.8.2+1693	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.8.2+312	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.8.2.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.8.2.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.8.2+176	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.8.2.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.8.2.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.8.2+14	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.8.2.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.8.2.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.8.2+86	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.8.2.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.8.2.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.8.2+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.8.2.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.8.2.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.8.2+250	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.8.2.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.8.2.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.8.2+434	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.8.2.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.8.2.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.8.2+218	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.8.2.tar.gz	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.8.2.html	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.8.2.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.8.2+74	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.8.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.8.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.8.2+40	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.8.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.8.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.8.2+19	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.8.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.8.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.8.2.html#changes
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.8.2+95	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.8.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.8.2.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.8.2.html#changes
		Tarball	Release notes	Changes

CDH 5.8.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.8.0+112	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.8.0+91	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.8.0+18	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.8.0+50	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.8.0.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/finag160cdh580.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/finag160cdh580.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/finag160cdh580CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.8.0+1592	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/haip260cdh580.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/haip260cdh580.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/haip260cdh580CHANGES.txt
		Tarball	Release notes	Changes
Hbase	hbase-1.2.0+cdh5.8.0+160	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbae120cdh580.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbae120cdh580.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbae120cdh580CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.8.0+64	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbas15cdh580.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbas15cdh580.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbas15cdh580CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.8.0+610	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hie110cdh580.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hie110cdh580.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hie110cdh580CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.8.0+2512	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/he390cdh580.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/he390cdh580.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/he390cdh580CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.6.0+cdh5.8.0+0	(none)	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/imp260cdh580.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/imp260cdh580CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.8.0+136	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/kite100cdh580.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/kite100cdh580.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/kite100cdh580CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.8.0			

Component	Package Version	Tarball	Release Notes	Changes File
	0+0	https://archive.apache.org/dist/parquet/parquet-1.0.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.0.0+cdh5.8.0.html	https://archive.apache.org/dist/parquet/parquet-1.0.0+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.8.0+27	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.8.0.tar.gz	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.8.0.html	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.8.0+291	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.8.0.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.8.0+174	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.8.0.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.8.0+12	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.8.0.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.8.0+83	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.8.0.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.8.0+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.8.0.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.8.0.html#changes
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.8.0+244	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.8.0.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.8.0.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.8.0.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Solr	solr-4.10.3+cdh5.8.0+423	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.10.3+cdh5.8.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.8.0+205	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.6.0+cdh5.8.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.8.0+65	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.8.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.8.0+38	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.8.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.8.0+17	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.8.0.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.8.0+94	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.8.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.8.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.8.0.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.7.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.7.x release are listed in the following table. For the same type of information for other CDH releases, see [Packaging and Tarballs](#). To view the overall release notes for CDH 5.x.x, see the [Release Notes](#).



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.7.6 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.7.6+119	https://archive.apache.org/dist/avro/1.7.6/cdh5.7.6/avro-1.7.6+cdh5.7.6.tar.gz	https://archive.apache.org/dist/avro/1.7.6/cdh5.7.6/avro-1.7.6+cdh5.7.6-releasenotes.html	https://archive.apache.org/dist/avro/1.7.6/cdh5.7.6/avro-1.7.6+cdh5.7.6-changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.7.6+97	https://archive.apache.org/dist/crunch/0.11.0/cdh5.7.6/crunch-0.11.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/crunch/0.11.0/cdh5.7.6/crunch-0.11.0+cdh5.7.6-releasenotes.html	https://archive.apache.org/dist/crunch/0.11.0/cdh5.7.6/crunch-0.11.0+cdh5.7.6-changes.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.7.6+24	https://archive.apache.org/dist/datafu/1.1.0/cdh5.7.6/datafu-1.1.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/datafu/1.1.0/cdh5.7.6/datafu-1.1.0+cdh5.7.6-releasenotes.html	https://archive.apache.org/dist/datafu/1.1.0/cdh5.7.6/datafu-1.1.0+cdh5.7.6-changes.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.7.6+61	https://archive.apache.org/dist/flume-ng/1.6.0/cdh5.7.6/flume-ng-1.6.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/flume-ng/1.6.0/cdh5.7.6/flume-ng-1.6.0+cdh5.7.6-releasenotes.html	https://archive.apache.org/dist/flume-ng/1.6.0/cdh5.7.6/flume-ng-1.6.0+cdh5.7.6-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.7.6+1551	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5.7.6/hadoop-2.6.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5.7.6/hadoop-2.6.0+cdh5.7.6-releasenotes.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5.7.6/hadoop-2.6.0+cdh5.7.6-changes.txt
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.7.6+1551	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.7.6+208	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.6.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.6.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.7.6+70	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.7.6.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.7.6.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.7.6+685	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.6.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.6.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.7.6+1881	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.6.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.6.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.5.0+cdh5.7.6+0	(none)	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/impala-2.5.0+cdh5.7.6.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/impala-2.5.0+cdh5.7.6.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.7.6+140	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.6.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.6.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.7.6+0	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.6.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.6.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.7.6+33	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.7.6.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.7.6.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1-112557033321-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.7.6+306	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.7.6.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.7.6.PAUS.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.7.6+180	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.7.6.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.7.6.PAUS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.7.6+18	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.7.6.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.7.6.PAUS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.7.6+92	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.7.6.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.7.6.PAUS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.7.6+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.7.6.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.7.6.PAUS.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.7.6+214	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.7.6.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.7.6.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.7.6.PAUS.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.7.6+430	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.7.6.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.7.6.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.7.6.PAUS.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.7.6+240	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.6.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.6.0+cdh5.7.6.html	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.6.0+cdh5.7.6.PAUS.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.7.6+87	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.6.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.6.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.7.6+44	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.6.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.6.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.7.6+23	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.6.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.6.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.7.6+98	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.6.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.6.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.6.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.7.5 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.7.5+116	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.7.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.7.5.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.7.5+96	(none)	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.5.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.5.CHANGES.txt
			Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.7.5+22	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.5.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.7.5+51	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.5.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.5.CHANGES.txt

Component	Package Version	Tarball	Release Notes	Changes File
		<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh575.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh575.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh575.html#changes</p> <p>Changes</p>
Apache Hadoop	hadoop-2.6.0+cdh5.7.5+1488	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh575.tar.gz</p> <p>Tarball</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh575.html</p> <p>Release notes</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/160cdh575.html#changes</p> <p>Changes</p>
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.7.5+1488	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.7.5+194	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.5.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.5.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.7.5+68	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.7.5.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.7.5.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.7.5+646	(none)	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.5.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.5.CHANGES.txt
			Release notes	Changes
Hue	hue-3.9.0+cdh5.7.5+1878	(none)	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.5.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.5.CHANGES.txt
			Release notes	Changes
Apache Impala	impala-2.5.0+cdh5.7.5+0	(none)	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/impala-2.5.0+cdh5.7.5.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/impala-2.5.0+cdh5.7.5.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.7.5+138	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.5.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.5.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.7.5+0	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.5.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.5.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.7.5+31	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.7.5.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.7.5.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.7.5+298	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh575.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh575-releasenotes.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/oozie-410cdh575-CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.7.5+178	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh575.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh575-releasenotes.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-150cdh575-CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.7.5+16	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh575.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh575-releasenotes.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-210cdh575-CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.7.5+88	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh575.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh575-releasenotes.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/pig-0120cdh575-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.7.5+0	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh575.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh575-releasenotes.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/search-100cdh575-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.7.5+206	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh575.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh575-releasenotes.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/sentry-151cdh575-CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.7.5+417	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh575.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh575-releasenotes.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/solr-4103cdh575-CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.7.5+224	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh575.tar.gz	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh575-releasenotes.html	https://usapoc.blob.core.windows.net/p/cdh5/cdh/5/spark-160cdh575-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.7.5+71	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.7.5+42	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.7.5+21	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.7.5+96	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.5.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.7.4 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.7.4+116	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.7.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.7.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.7.4+95	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.7.4+22	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.7.4+51	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.7.4.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.7.4-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.7.4-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.7.4+1452	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.7.4.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.7.4-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.6.0-cdh5.7.4-changes.txt
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.7.4+1452	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.7.4+190	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.4.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.4.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.7.4+68	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.7.4.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.7.4.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.7.4+626	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.4.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.4.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.7.4+1863	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.4.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.4.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.5.0+cdh5.7.4+0	(none)	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/impala-2.5.0+cdh5.7.4.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/impala-2.5.0+cdh5.7.4.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.7.4+138	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.4.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.4.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.7.4+0	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.4.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.4.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.7.4+31	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.7.4.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.7.4.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.7.4.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.7.4+297	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/oozie-410cdh574.tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/oozie-410cdh574.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/oozie-410cdh574CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.7.4+178	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/parquet-150cdh574.tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/parquet-150cdh574.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/parquet-150cdh574CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.7.4+16	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/parquet-format-210cdh574.tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/parquet-format-210cdh574.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/parquet-format-210cdh574CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.7.4+87	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/pig-0120cdh574.tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/pig-0120cdh574.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/pig-0120cdh574CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.7.4+0	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/search-100cdh574.tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/search-100cdh574.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/search-100cdh574CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.7.4+205	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/sentry-151cdh574.tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/sentry-151cdh574.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/sentry-151cdh574CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.7.4+416	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/solr-4103cdh574.tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/solr-4103cdh574.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/solr-4103cdh574CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.7.4+211	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/spark-160cdh574.tar.gz	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/spark-160cdh574.html	https://usapocdn1.safelinks.com/p/cdh5/cdh/5/spark-160cdh574CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.7.4+68	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.4-release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.4-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.7.4+42	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.4-release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.4-CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.7.4+21	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.4-release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.4-CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.7.4+96	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.4-release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.4-CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.7.3 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.7.3+115	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.7.3-release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.7.3-CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.7.3+94	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.3-release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.3-CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.7.3+21	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.3-release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.3-CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.7.3+48	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.3-release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.3-CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		<p>p/cdh5/cdh/5/</p> <p>https://www.cloudera.com/p/cdh5/cdh/5/</p> <p>map-260-cdh-7.3-1408.tgz</p>	<p>p/cdh5/cdh/5/</p> <p>https://www.cloudera.com/p/cdh5/cdh/5/</p> <p>map-260-cdh-7.3-1408.tgz</p>	<p>p/cdh5/cdh/5/</p> <p>https://www.cloudera.com/p/cdh5/cdh/5/</p> <p>map-260-cdh-7.3-1408.tgz</p>
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.7.3+1408	(none)	<p>https://www.cloudera.com/p/cdh5/cdh/5/</p> <p>https://www.cloudera.com/p/cdh5/cdh/5/</p> <p>map-260-cdh-7.3-1408.tgz</p>	<p>https://www.cloudera.com/p/cdh5/cdh/5/</p> <p>https://www.cloudera.com/p/cdh5/cdh/5/</p> <p>map-260-cdh-7.3-1408.tgz</p>
			Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.7.3+1408	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.2.0+cdh5.7.3+176	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.3.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.3.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.7.3+67	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.7.3.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.7.3.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hbase-1.5+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.7.3+607	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.3.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.3.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.7.3+1855	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.3.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.3.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.5.0+cdh5.7.3+0	(none)	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/impala-2.5.0+cdh5.7.3.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/impala-2.5.0+cdh5.7.3.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.7.3+137	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.3.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.3.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.7.3+0	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.3.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.3.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.7.3+30	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.7.3.tar.gz	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.7.3.release.html	https://s3.amazonaws.com/elasticmapreduce-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.7.3+285	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.7.3.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.7.3.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.7.3.PM5.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.7.3+177	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.7.3.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.7.3.PM5.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.7.3+15	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.7.3.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.7.3.PM5.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.7.3+85	(none)	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.7.3.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.7.3.PM5.txt
			Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.7.3+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.7.3.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.7.3.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.7.3.PM5.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.7.3+204	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.7.3.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.7.3.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.7.3.PM5.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.7.3+414	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.7.3.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.7.3.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.7.3.PM5.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.7.3+206	(none)	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.7.3.html	https://archive.apache.org/dist/spark/spark-1.6.0+cdh5.7.3.PM5.txt
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.7.3+67	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.3.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.7.3+41	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.3.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.7.3+20	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.3.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.7.3+95	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.3.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.3.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.7.3.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.7.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.7.2+114	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.7.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.7.2.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.7.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.7.2+93	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.2.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.7.2.CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.7.2+20	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.2.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.7.2.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.7.2+44	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.2.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.2.release.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.2.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/finag160cdh572.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/finag160cdh572.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/finag160cdh572CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.7.2+1380	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hadoop260cdh572.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hadoop260cdh572.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hadoop260cdh572CHANGES.txt
		Tarball	Release notes	Changes
Hbase	hbase-1.2.0+cdh5.7.2+166	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase120cdh572.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase120cdh572.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase120cdh572CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.7.2+65	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase15cdh572.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase15cdh572.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase15cdh572CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.7.2+587	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive110cdh572.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive110cdh572.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive110cdh572CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.7.2+1841	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue390cdh572.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue390cdh572.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue390cdh572CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.5.0+cdh5.7.2+0	(none)	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/impala250cdh572.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/impala250cdh572CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.7.2+136	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite100cdh572.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite100cdh572.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite100cdh572CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.7.2+0			

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
	2+0	https://archive.apache.org/dist/parquet/parquet-1.0.0+cdh5.7.2.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.0.0+cdh5.7.2-release.html	https://archive.apache.org/dist/parquet/parquet-1.0.0+cdh5.7.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.7.2+29	(none)	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.7.2-release.html	https://archive.apache.org/dist/mahout/mahout-0.9+cdh5.7.2-CHANGES.txt
			Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.7.2+282	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.7.2.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.7.2-release.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.7.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.7.2+176	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.7.2.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.7.2-release.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.7.2-CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.7.2+14	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.7.2.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.7.2-release.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.7.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.7.2+84	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.7.2.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.7.2-release.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.7.2-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.7.2+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.7.2.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.7.2-release.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.7.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.7.2+203	(none)	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.7.2-release.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.7.2-CHANGES.txt
			Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Solr	solr-4.10.3+cdh5.7.2+410	(none)	https://archive.apache.org/dist/cdh/5/solr-4.10.3+cdh5.7.2+410/	https://archive.apache.org/dist/cdh/5/solr-4.10.3+cdh5.7.2+410/changes.txt
			Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.7.2+201	https://archive.apache.org/dist/cdh/5/spark-1.6.0+cdh5.7.2+201.tar.gz	https://archive.apache.org/dist/cdh/5/spark-1.6.0+cdh5.7.2+201/	https://archive.apache.org/dist/cdh/5/spark-1.6.0+cdh5.7.2+201/changes.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.7.2+62	(none)	https://archive.apache.org/dist/cdh/5/sqoop-1.4.6+cdh5.7.2+62/	https://archive.apache.org/dist/cdh/5/sqoop-1.4.6+cdh5.7.2+62/changes.txt
			Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.7.2+40	https://archive.apache.org/dist/cdh/5/sqoop2-1.99.5+cdh5.7.2+40.tar.gz	https://archive.apache.org/dist/cdh/5/sqoop2-1.99.5+cdh5.7.2+40/	https://archive.apache.org/dist/cdh/5/sqoop2-1.99.5+cdh5.7.2+40/changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.7.2+19	https://archive.apache.org/dist/cdh/5/whirr-0.9.0+cdh5.7.2+19.tar.gz	https://archive.apache.org/dist/cdh/5/whirr-0.9.0+cdh5.7.2+19/	https://archive.apache.org/dist/cdh/5/whirr-0.9.0+cdh5.7.2+19/changes.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.7.2+94	https://archive.apache.org/dist/cdh/5/zookeeper-3.4.5+cdh5.7.2+94.tar.gz	https://archive.apache.org/dist/cdh/5/zookeeper-3.4.5+cdh5.7.2+94/	https://archive.apache.org/dist/cdh/5/zookeeper-3.4.5+cdh5.7.2+94/changes.txt
		Tarball	Release notes	Changes

CDH 5.7.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.7.1+113	https://archive.apache.org/dist/cdh/5/avro-1.7.6+cdh5.7.1+113.tar.gz	https://archive.apache.org/dist/cdh/5/avro-1.7.6+cdh5.7.1+113/	https://archive.apache.org/dist/cdh/5/avro-1.7.6+cdh5.7.1+113/changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.7.1+92	https://archive.apache.org/dist/cdh/5/crunch-0.11.0+cdh5.7.1+92.tar.gz	https://archive.apache.org/dist/cdh/5/crunch-0.11.0+cdh5.7.1+92/	https://archive.apache.org/dist/cdh/5/crunch-0.11.0+cdh5.7.1+92/changes.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		p/cdh5/cdh/5/cdh5.0.110+cdh5.7.1.tar.gz	p/cdh5/cdh/5/cdh5.0.110+cdh5.7.1.tar.gz	p/cdh5/cdh/5/cdh5.0.110+cdh5.7.1.CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.7.1+19	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/datafu-1.1.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/datafu-1.1.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/datafu-1.1.0+cdh5.7.1.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.7.1+40	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.7.1+1335	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.7.1.CHANGES.txt
		Tarball	Release notes	Changes
Hbase	hbase-1.2.0+cdh5.7.1+142	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-1.2.0+cdh5.7.1.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.7.1+64	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.7.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.7.1+564	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-1.1.0+cdh5.7.1.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.7.1+1804	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.1.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.9.0+cdh5.7.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.5.0+cdh5.7.	(none)		

Component	Package Version	Tarball	Release Notes	Changes File
	1+0		https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ parquet-150cdh571.html	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ parquet-150cdh571.html
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.7.1+135	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ kite-100cdh571.tar.gz	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ kite-100cdh571.html	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ kite-100cdh571.html
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.7.1+0	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ llama-100cdh571.tar.gz	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ llama-100cdh571.html	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ llama-100cdh571.html
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.7.1+28	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ mahout-09cdh571.tar.gz	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ mahout-09cdh571.html	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ mahout-09cdh571.html
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.7.1+271	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ oozie-410cdh571.tar.gz	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ oozie-410cdh571.html	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ oozie-410cdh571.html
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.7.1+175	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ parquet-150cdh571.tar.gz	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ parquet-150cdh571.html	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ parquet-150cdh571.html
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.7.1+13	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ parquet-format-210cdh571.tar.gz	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ parquet-format-210cdh571.html	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ parquet-format-210cdh571.html
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.7.1+83	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ pig-0120cdh571.tar.gz	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ pig-0120cdh571.html	https:// userguide.cloudera.com/ p/cdh5/cdh/ 5/ pig-0120cdh571.html
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Cloudera Search	search-1.0.0+cdh5.7.1+0	https://archive.cloudera.com/p/cdh5/cdh/5/search100cdh571.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/search100cdh571.html	https://archive.cloudera.com/p/cdh5/cdh/5/search100cdh571Changes.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.7.1+190	https://archive.cloudera.com/p/cdh5/cdh/5/sentry151cdh571.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/sentry151cdh571.html	https://archive.cloudera.com/p/cdh5/cdh/5/sentry151cdh571Changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.7.1+403	https://archive.cloudera.com/p/cdh5/cdh/5/solr4103cdh571.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/solr4103cdh571.html	https://archive.cloudera.com/p/cdh5/cdh/5/solr4103cdh571Changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.7.1+193	https://archive.cloudera.com/p/cdh5/cdh/5/spark160cdh571.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/spark160cdh571.html	https://archive.cloudera.com/p/cdh5/cdh/5/spark160cdh571Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.7.1+56	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop146cdh571.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop146cdh571.html	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop146cdh571Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.7.1+39	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop21995cdh571.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop21995cdh571.html	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop21995cdh571Changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.7.1+18	https://archive.cloudera.com/p/cdh5/cdh/5/whirr090cdh571.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/whirr090cdh571.html	https://archive.cloudera.com/p/cdh5/cdh/5/whirr090cdh571Changes.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.7.1+93	https://archive.cloudera.com/p/cdh5/cdh/5/zookeeper345cdh571.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/zookeeper345cdh571.html	https://archive.cloudera.com/p/cdh5/cdh/5/zookeeper345cdh571Changes.txt
		Tarball	Release notes	Changes

CDH 5.7.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.7.0+113	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.7.0.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.7.0.release.html	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.7.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.7.0+93	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.7.0.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.7.0.release.html	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.7.0.CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.7.0+20	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.10+cdh5.7.0.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.10+cdh5.7.0.release.html	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.10+cdh5.7.0.CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.7.0+37	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.0.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.0.release.html	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.6.0+cdh5.7.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.7.0+1280	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.7.0.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.7.0.release.html	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.7.0.CHANGES.txt
		Tarball	Release notes	Changes
HBase	hbase-1.2.0+cdh5.7.0+129	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.2.0+cdh5.7.0.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.2.0+cdh5.7.0.release.html	https://us.apache.org/dist/cdh5/cdh/5/hbase-1.2.0+cdh5.7.0.CHANGES.txt
		Tarball	Release notes	Changes
HBase-solr	hbase-solr-1.5+cdh5.7.0+64	https://us.apache.org/dist/cdh5/cdh/5/hbase-solr-1.5+cdh5.7.0.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hbase-solr-1.5+cdh5.7.0.release.html	https://us.apache.org/dist/cdh5/cdh/5/hbase-solr-1.5+cdh5.7.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.7.0+522	https://us.apache.org/dist/cdh5/cdh/	https://us.apache.org/dist/cdh5/cdh/	https://us.apache.org/dist/cdh5/cdh/

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		5/ hive-1.10.0-cdh5.7.0.tar.gz	5/ hive-1.10.0-cdh5.7.0-release.html	5/ hive-1.10.0-cdh5.7.0-changes.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.7.0+1759	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ hue-3.9.0-cdh5.7.0.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ hue-3.9.0-cdh5.7.0-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ hue-3.9.0-cdh5.7.0-changes.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.5.0+cdh5.7.0+0	(none)	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ impala-2.5.0-cdh5.7.0-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ impala-2.5.0-cdh5.7.0-changes.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.7.0+130	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ kite-1.0.0-cdh5.7.0.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ kite-1.0.0-cdh5.7.0-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ kite-1.0.0-cdh5.7.0-changes.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.7.0+0	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.7.0.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.7.0-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.7.0-changes.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.7.0+29	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ mahout-0.9-cdh5.7.0.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ mahout-0.9-cdh5.7.0-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ mahout-0.9-cdh5.7.0-changes.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.7.0+267	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ oozie-4.1.0-cdh5.7.0.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ oozie-4.1.0-cdh5.7.0-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ oozie-4.1.0-cdh5.7.0-changes.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.7.0+176	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ parquet-1.5.0-cdh5.7.0.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ parquet-1.5.0-cdh5.7.0-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ parquet-1.5.0-cdh5.7.0-changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Parquet-format	parquet-format-2.10+cdh5.7.0+14	https://us.apache.org/dist/parquet/cdh5/cdh/5/parquet-format-2.10+cdh5.7.0.tar.gz	https://us.apache.org/dist/parquet/cdh5/cdh/5/parquet-format-2.10+cdh5.7.0.html	https://us.apache.org/dist/parquet/cdh5/cdh/5/parquet-format-2.10+cdh5.7.0.PAGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.7.0+84	https://us.apache.org/dist/pig/cdh5/cdh/5/pig-0.12.0+cdh5.7.0.tar.gz	https://us.apache.org/dist/pig/cdh5/cdh/5/pig-0.12.0+cdh5.7.0.html	https://us.apache.org/dist/pig/cdh5/cdh/5/pig-0.12.0+cdh5.7.0.PAGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.7.0+0	https://us.apache.org/dist/search/cdh5/cdh/5/search-1.0.0+cdh5.7.0.tar.gz	https://us.apache.org/dist/search/cdh5/cdh/5/search-1.0.0+cdh5.7.0.html	https://us.apache.org/dist/search/cdh5/cdh/5/search-1.0.0+cdh5.7.0.PAGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.7.0+184	https://us.apache.org/dist/sentry/cdh5/cdh/5/sentry-1.5.1+cdh5.7.0.tar.gz	https://us.apache.org/dist/sentry/cdh5/cdh/5/sentry-1.5.1+cdh5.7.0.html	https://us.apache.org/dist/sentry/cdh5/cdh/5/sentry-1.5.1+cdh5.7.0.PAGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.7.0+389	https://us.apache.org/dist/solr/cdh5/cdh/5/solr-4.10.3+cdh5.7.0.tar.gz	https://us.apache.org/dist/solr/cdh5/cdh/5/solr-4.10.3+cdh5.7.0.html	https://us.apache.org/dist/solr/cdh5/cdh/5/solr-4.10.3+cdh5.7.0.PAGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.6.0+cdh5.7.0+180	https://us.apache.org/dist/spark/cdh5/cdh/5/spark-1.6.0+cdh5.7.0.tar.gz	https://us.apache.org/dist/spark/cdh5/cdh/5/spark-1.6.0+cdh5.7.0.html	https://us.apache.org/dist/spark/cdh5/cdh/5/spark-1.6.0+cdh5.7.0.PAGES.txt
		Tarball	Release notes	Changes
Spark-netlib	spark-netlib-master.2	(none)	https://us.apache.org/dist/spark-netlib/cdh5/cdh/5/spark-netlib-master.2.html	https://us.apache.org/dist/spark-netlib/cdh5/cdh/5/spark-netlib-master.2.PAGES.txt
			Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.7.0+56	https://us.apache.org/dist/sqoop/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.0.tar.gz	https://us.apache.org/dist/sqoop/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.0.html	https://us.apache.org/dist/sqoop/cdh5/cdh/5/sqoop-1.4.6+cdh5.7.0.PAGES.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.7.0+38	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.0.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.0.html	https://archive.cloudera.com/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.7.0.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.7.0+19	https://archive.cloudera.com/p/cdh5/cdh/5/whirr-0.9.0+cdh5.7.0.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/whirr-0.9.0+cdh5.7.0.html	https://archive.cloudera.com/p/cdh5/cdh/5/whirr-0.9.0+cdh5.7.0.html#changes
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.7.0+94	https://archive.cloudera.com/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.7.0.tar.gz	https://archive.cloudera.com/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.7.0.html	https://archive.cloudera.com/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.7.0.html#changes
		Tarball	Release notes	Changes

CDH 5.6.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.6.x release are listed in the following table. For the same type of information for other CDH releases, see [Packaging and Tarballs](#). To view the overall release notes for CDH 5.x.x, see the [Release Notes](#).



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.6.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.6.1+113	https://archive.apache.org/dist/avro/1.7.6/avro-1.7.6-cdh5.6.1.tar.gz	https://archive.apache.org/dist/avro/1.7.6/avro-1.7.6-cdh5.6.1-releasenotes.html	https://archive.apache.org/dist/avro/1.7.6/avro-1.7.6-cdh5.6.1-changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.6.1+78	https://archive.apache.org/dist/crunch/0.11.0/crunch-0.11.0-cdh5.6.1.tar.gz	https://archive.apache.org/dist/crunch/0.11.0/crunch-0.11.0-cdh5.6.1-releasenotes.html	https://archive.apache.org/dist/crunch/0.11.0/crunch-0.11.0-cdh5.6.1-changes.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.6.1+18	https://archive.apache.org/dist/datafu/1.1.0/datafu-1.1.0-cdh5.6.1.tar.gz	https://archive.apache.org/dist/datafu/1.1.0/datafu-1.1.0-cdh5.6.1-releasenotes.html	https://archive.apache.org/dist/datafu/1.1.0/datafu-1.1.0-cdh5.6.1-changes.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.6.1+38	https://archive.apache.org/dist/flume-ng/1.6.0/flume-ng-1.6.0-cdh5.6.1.tar.gz	https://archive.apache.org/dist/flume-ng/1.6.0/flume-ng-1.6.0-cdh5.6.1-releasenotes.html	https://archive.apache.org/dist/flume-ng/1.6.0/flume-ng-1.6.0-cdh5.6.1-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.6.1+1092	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.6.1.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.6.1-releasenotes.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.6.1-changes.txt
		Tarball	Release notes	Changes
HBase	hbase-1.0.0+cdh5.6.1+314	https://archive.apache.org/dist/hbase/1.0.0/hbase-1.0.0-cdh5.6.1.tar.gz	https://archive.apache.org/dist/hbase/1.0.0/hbase-1.0.0-cdh5.6.1-releasenotes.html	https://archive.apache.org/dist/hbase/1.0.0/hbase-1.0.0-cdh5.6.1-changes.txt
		Tarball	Release notes	Changes
HBase-solr	hbase-solr-1.5+cdh5.6.1+59	https://archive.apache.org/dist/hbase-solr/1.5/hbase-solr-1.5-cdh5.6.1.tar.gz	https://archive.apache.org/dist/hbase-solr/1.5/hbase-solr-1.5-cdh5.6.1-releasenotes.html	https://archive.apache.org/dist/hbase-solr/1.5/hbase-solr-1.5-cdh5.6.1-changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.6.1+429	https://archive.apache.org/dist/hive/1.1.0/hive-1.1.0-cdh5.6.1.tar.gz	https://archive.apache.org/dist/hive/1.1.0/hive-1.1.0-cdh5.6.1-releasenotes.html	https://archive.apache.org/dist/hive/1.1.0/hive-1.1.0-cdh5.6.1-changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.6.1+391	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue390cdh561.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue390cdh561.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue390cdh561CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.4.0+cdh5.6.1+0	(none)	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/impala240cdh561.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/impala240cdh561CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.6.1+123	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite100cdh561.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite100cdh561.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite100cdh561CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.6.1+0	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama100cdh561.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama100cdh561.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama100cdh561CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.6.1+27	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/mahout09cdh561.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/mahout09cdh561.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/mahout09cdh561CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.6.1+242	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/oozie410cdh561.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/oozie410cdh561.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/oozie410cdh561CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.6.1+175	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/parquet150cdh561.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/parquet150cdh561.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/parquet150cdh561CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.6.1+13	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/parquet-format210cdh561.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/parquet-format210cdh561.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/parquet-format210cdh561CHANGES.txt

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.6.1+74	https://archive.apache.org/dist/parcels/cdh5/cdh/5/pig0120cdh561.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/pig0120cdh561.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/pig0120cdh561Changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.6.1+0	https://archive.apache.org/dist/parcels/cdh5/cdh/5/search100cdh561.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/search100cdh561.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/search100cdh561Changes.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.6.1+137	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry151cdh561.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry151cdh561.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry151cdh561Changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.6.1+368	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr4103cdh561.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr4103cdh561.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr4103cdh561Changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.5.0+cdh5.6.1+130	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark150cdh561.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark150cdh561.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark150cdh561Changes.txt
		Tarball	Release notes	Changes
Spark-netlib	spark-netlib-master.1	(none)	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sparknetlib1cdh561.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sparknetlib1cdh561Changes.txt
			Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.6.1+36	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop146cdh561.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop146cdh561.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop146cdh561Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.6.1+37	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop21995cdh561.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop21995cdh561.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop21995cdh561Changes.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.6.1+18	https://archive.apache.org/dist/whirr/p/cdh5/cdh/5/whirr-0.9.0+cdh5.6.1.tar.gz	https://archive.apache.org/dist/whirr/p/cdh5/cdh/5/whirr-0.9.0+cdh5.6.1.html	https://archive.apache.org/dist/whirr/p/cdh5/cdh/5/whirr-0.9.0+cdh5.6.1.html
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.6.1+92	https://archive.apache.org/dist/zookeeper/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.6.1.tar.gz	https://archive.apache.org/dist/zookeeper/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.6.1.html	https://archive.apache.org/dist/zookeeper/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.6.1.html
		Tarball	Release notes	Changes

CDH 5.6.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.6.0+112	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.6+cdh5.6.0.tar.gz	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.6+cdh5.6.0.html	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.6+cdh5.6.0.html
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.6.0+77	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.11.0+cdh5.6.0.tar.gz	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.11.0+cdh5.6.0.html	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.11.0+cdh5.6.0.html
		Tarball	Release notes	Changes
Datafu	pig-df-datafu-1.10+cdh5.6.0+17	https://archive.apache.org/dist/datafu/p/cdh5/cdh/5/datafu-1.10+cdh5.6.0.tar.gz	https://archive.apache.org/dist/datafu/p/cdh5/cdh/5/datafu-1.10+cdh5.6.0.html	https://archive.apache.org/dist/datafu/p/cdh5/cdh/5/datafu-1.10+cdh5.6.0.html
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.6.0+30	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.6.0.tar.gz	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.6.0.html	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.6.0+cdh5.6.0.html
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.6.0+1025	https://archive.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.6.0.tar.gz	https://archive.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.6.0.html	https://archive.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.6.0.html
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Hadoop Mrv1	http://mrvc.2015.10.15	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.0.0+cdh5.6.0+300	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.0.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.0.0+cdh5.6.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.0.0+cdh5.6.0.PM5.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.6.0+57	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.6.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.6.0.PM5.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.6.0+377	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.6.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.6.0.PM5.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.6.0+365	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.6.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.6.0.PM5.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.4.0+cdh5.6.0+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.4.0+cdh5.6.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.4.0+cdh5.6.0.PM5.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.6.0+119	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.6.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.6.0.PM5.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.6.0+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.6.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.6.0.PM5.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.6.0+26	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.6.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.6.0.PM5.txt

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.6.0+235	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.1.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.1.0+cdh5.6.0.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.1.0+cdh5.6.0.PMS.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.6.0+174	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.6.0.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.6.0.PMS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.6.0+12	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.6.0.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.6.0.PMS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.6.0+73	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.6.0.html	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.6.0.PMS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.6.0+0	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.6.0.html	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.6.0.PMS.txt
		Tarball	Release notes	Changes
Apache Sentry (Incubating)	sentry-1.5.1+cdh5.6.0+121	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.5.1+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.5.1+cdh5.6.0.html	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.5.1+cdh5.6.0.PMS.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.6.0+339	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.10.3+cdh5.6.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.10.3+cdh5.6.0.html	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.10.3+cdh5.6.0.PMS.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.5.0+cdh5.6.0+115	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		5/ parquet-1.5.0-cdh5.6.0.tar.gz	5/ parquet-1.5.0-cdh5.6.0-release.html	5/ parquet-1.5.0-cdh5.6.0-CHANGES.txt
		Tarball	Release notes	Changes
Spark-netlib	spark-netlib-master.2	(none)	https://archive.apache.org/dist/spark/spark-netlib-master.2-cdh5/cdh/5/spark-netlib-master.2-cdh5.6.0-release.html	https://archive.apache.org/dist/spark/spark-netlib-master.2-cdh5/cdh/5/spark-netlib-master.2-cdh5.6.0-CHANGES.txt
			Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.6.0+33	https://archive.apache.org/dist/sqoop/sqoop-1.4.6-cdh5.6.0.tar.gz	https://archive.apache.org/dist/sqoop/sqoop-1.4.6-cdh5.6.0-release.html	https://archive.apache.org/dist/sqoop/sqoop-1.4.6-cdh5.6.0-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.6.0+34	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.5-cdh5.6.0.tar.gz	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.5-cdh5.6.0-release.html	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.5-cdh5.6.0-CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.6.0+17	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.6.0.tar.gz	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.6.0-release.html	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.6.0-CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.6.0+91	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.6.0.tar.gz	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.6.0-release.html	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.6.0-CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.5.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.5.x release are listed in the following table. For the same type of information for other CDH releases, see [Packaging and Tarballs](#). To view the overall release notes for CDH 5.x.x, see the [Release Notes](#).



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.5.6 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.5.6+120	https://archive.apache.org/dist/avro/1.7.6/cdh5.5.6/avro-1.7.6-cdh5.5.6.tar.gz	https://archive.apache.org/dist/avro/1.7.6/cdh5.5.6/avro-1.7.6-cdh5.5.6.html	https://archive.apache.org/dist/avro/1.7.6/cdh5.5.6/avro-1.7.6-cdh5.5.6.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.5.6+85	https://archive.apache.org/dist/crunch/0.11.0/cdh5.5.6/crunch-0.11.0-cdh5.5.6.tar.gz	https://archive.apache.org/dist/crunch/0.11.0/cdh5.5.6/crunch-0.11.0-cdh5.5.6.html	https://archive.apache.org/dist/crunch/0.11.0/cdh5.5.6/crunch-0.11.0-cdh5.5.6.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.5.6+25	https://archive.apache.org/dist/datafu/1.1.0/cdh5.5.6/datafu-1.1.0-cdh5.5.6.tar.gz	https://archive.apache.org/dist/datafu/1.1.0/cdh5.5.6/datafu-1.1.0-cdh5.5.6.html	https://archive.apache.org/dist/datafu/1.1.0/cdh5.5.6/datafu-1.1.0-cdh5.5.6.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.5.6+51	https://archive.apache.org/dist/flume-ng/1.6.0/cdh5.5.6/flume-ng-1.6.0-cdh5.5.6.tar.gz	https://archive.apache.org/dist/flume-ng/1.6.0/cdh5.5.6/flume-ng-1.6.0-cdh5.5.6.html	https://archive.apache.org/dist/flume-ng/1.6.0/cdh5.5.6/flume-ng-1.6.0-cdh5.5.6.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.5.6+1242	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.5.6.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.5.6.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.5.6.txt
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.5.6+1242	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.0.0+cdh5.5.6+354	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.0.0+cdh5.5.6.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.0.0+cdh5.5.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.0.0+cdh5.5.6.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.5.6+66	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.5.6.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.5.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-1.5+cdh5.5.6.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.5.6+477	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.5.6.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.5.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-1.1.0+cdh5.5.6.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.5.6+426	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.5.6.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.5.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.9.0+cdh5.5.6.CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.3.0+cdh5.5.6+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.3.0+cdh5.5.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.3.0+cdh5.5.6.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.5.6+130	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.5.6.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.5.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-1.0.0+cdh5.5.6.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.5.6+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.5.6.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.5.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.5.6.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.5.6+34	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.5.6.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.5.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.5.6.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.5.6+273	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.5.6.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.5.6.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.5.6/changes.html
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.5.6+182	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.5.6.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.5.6.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.5.6/changes.html
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.5.6+20	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.5.6.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.5.6.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.5.6/changes.html
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.5.6+85	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.5.6.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.5.6.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.5.6/changes.html
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.5.6+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.5.6.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.5.6.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.5.6/changes.html
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.5.6+150	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.5.6.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.5.6.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.5.6/changes.html
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.5.6+391	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.5.6.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.5.6.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.5.6/changes.html
		Tarball	Release notes	Changes
Apache Spark	spark-1.5.0+cdh5.5.6+144	https://archive.apache.org/dist/spark/spark-1.5.0+cdh5.5.6.tar.gz	https://archive.apache.org/dist/spark/spark-1.5.0+cdh5.5.6.html	https://archive.apache.org/dist/spark/spark-1.5.0+cdh5.5.6/changes.html
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.5.6+58	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.5.6.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.5.6.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.5.6.changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.5.6+44	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.5.6.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.5.6.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.5.6.changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.5.6+25	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.5.6.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.5.6.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.5.6.changes.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.5.6+99	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.5.6.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.5.6.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.5.6.changes.txt
		Tarball	Release notes	Changes

CDH 5.5.5 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.5.5+119	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.5.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.5.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.5.5.changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.5.5+84	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.5.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.5.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.5.5.changes.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.5.5+24	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.5.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.5.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.1.0+cdh5.5.5.changes.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.5.5+47	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.5.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.5.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.5.5.changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		https://www.cloudera.com/p/cdh5/cdh/5/160cdh555.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/160cdh555.html	https://www.cloudera.com/p/cdh5/cdh/5/160cdh555CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.5.5+1140	https://www.cloudera.com/p/cdh5/cdh/5/160cdh555.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/160cdh555.html	https://www.cloudera.com/p/cdh5/cdh/5/160cdh555CHANGES.txt
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.5.5+1140	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.0.0+cdh5.5+335	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hbase-1.0.0+cdh5.5.tar.gz	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hbase-1.0.0+cdh5.5.release.html	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hbase-1.0.0+cdh5.5.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.5.5+65	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.5.5.tar.gz	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.5.5.release.html	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.5.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.5+449	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.5.tar.gz	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.5.release.html	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hive-1.1.0+cdh5.5.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.5+421	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.5.tar.gz	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.5.release.html	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/hue-3.9.0+cdh5.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.3.0+cdh5.5+0	(none)	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/impala-2.3.0+cdh5.5.release.html	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/impala-2.3.0+cdh5.5.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.5+129	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.5.tar.gz	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.5.release.html	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/kite-1.0.0+cdh5.5.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.5+0	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.5.tar.gz	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.5.release.html	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/llama-1.0.0+cdh5.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.5+33	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.5.tar.gz	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.5.release.html	https://s3.amazonaws.com/elasticmaven-us-east-1-11255703333-us-east-1/p/cdh5/cdh/5/mahout-0.9+cdh5.5.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.5+260	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.5.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.5.html	https://archive.apache.org/dist/oozie/oozie-4.1.0+cdh5.5/changes.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.5+181	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.5.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.5.html	https://archive.apache.org/dist/parquet/parquet-1.5.0+cdh5.5/changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.5+19	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.5.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.5.html	https://archive.apache.org/dist/parquet-format/parquet-format-2.1.0+cdh5.5/changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.5+81	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.5.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.5.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.5/changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.5+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.5.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.5.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.5/changes.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.5+148	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.5.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.5.html	https://archive.apache.org/dist/sentry/sentry-1.5.1+cdh5.5/changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.5+380	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.5.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.5.html	https://archive.apache.org/dist/solr/solr-4.10.3+cdh5.5/changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.5.0+cdh5.5+140	https://archive.apache.org/dist/spark/spark-1.5.0+cdh5.5.tar.gz	https://archive.apache.org/dist/spark/spark-1.5.0+cdh5.5.html	https://archive.apache.org/dist/spark/spark-1.5.0+cdh5.5/changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.6+cdh5.5.5+43	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.5.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.5.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.6+cdh5.5.5.html
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.5.5+43	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.5.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.5.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.5.5.html
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.5.5+24	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.5.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.5.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.5.5.html
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.5.5+98	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.5.5.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.5.5.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.5.5.html
		Tarball	Release notes	Changes

CDH 5.5.4 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.5.4+118	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.5.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.5.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/avro-1.7.6+cdh5.5.4.html
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.5.4+83	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.5.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.5.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.5.4.html
		Tarball	Release notes	Changes
Datafu	pigudatafu-1.10+cdh5.5.4+23	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.10+cdh5.5.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.10+cdh5.5.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/datafu-1.10+cdh5.5.4.html
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.6.0+cdh5.5.4+43	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.5.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.5.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.6.0+cdh5.5.4.html
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.5.4.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.5.4-releasenotes.html	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.5.4-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.5.4+1072	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.5.4.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.5.4-releasenotes.html	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.5.4-changes.txt
		Tarball	Release notes	Changes
HBase	hbase-1.0.0+cdh5.5.4+318	https://us.apache.org/dist/hadoop/hbase/hbase-1.0.0+cdh5.5.4.tar.gz	https://us.apache.org/dist/hadoop/hbase/hbase-1.0.0+cdh5.5.4-releasenotes.html	https://us.apache.org/dist/hadoop/hbase/hbase-1.0.0+cdh5.5.4-changes.txt
		Tarball	Release notes	Changes
HBase-solr	hbase-solr-1.5+cdh5.5.4+64	https://us.apache.org/dist/hadoop/hbase/hbase-solr-1.5+cdh5.5.4.tar.gz	https://us.apache.org/dist/hadoop/hbase/hbase-solr-1.5+cdh5.5.4-releasenotes.html	https://us.apache.org/dist/hadoop/hbase/hbase-solr-1.5+cdh5.5.4-changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.5.4+422	https://us.apache.org/dist/hadoop/hive/hive-1.1.0+cdh5.5.4.tar.gz	https://us.apache.org/dist/hadoop/hive/hive-1.1.0+cdh5.5.4-releasenotes.html	https://us.apache.org/dist/hadoop/hive/hive-1.1.0+cdh5.5.4-changes.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.5.4+391	https://us.apache.org/dist/hadoop/hue/hue-3.9.0+cdh5.5.4.tar.gz	https://us.apache.org/dist/hadoop/hue/hue-3.9.0+cdh5.5.4-releasenotes.html	https://us.apache.org/dist/hadoop/hue/hue-3.9.0+cdh5.5.4-changes.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.3.0+cdh5.5.4+0	(none)	https://us.apache.org/dist/hadoop/impala/impala-2.3.0+cdh5.5.4-releasenotes.html	https://us.apache.org/dist/hadoop/impala/impala-2.3.0+cdh5.5.4-changes.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.5.4+128	https://us.apache.org/dist/hadoop/kite/kite-1.0.0+cdh5.5.4.tar.gz	https://us.apache.org/dist/hadoop/kite/kite-1.0.0+cdh5.5.4-releasenotes.html	https://us.apache.org/dist/hadoop/kite/kite-1.0.0+cdh5.5.4-changes.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.5.4+0	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.tar.gz	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.html	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.PMS.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.5.4+32	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/mahout-0.9-cdh5.4.tar.gz	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/mahout-0.9-cdh5.4.html	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/mahout-0.9-cdh5.4.PMS.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.5.4+247	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.4.tar.gz	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.4.html	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.4.PMS.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.5.4+180	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/parquet-1.5.0-cdh5.4.tar.gz	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/parquet-1.5.0-cdh5.4.html	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/parquet-1.5.0-cdh5.4.PMS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.5.4+18	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.4.tar.gz	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.4.html	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.4.PMS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.5.4+80	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/pig-0.12.0-cdh5.4.tar.gz	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/pig-0.12.0-cdh5.4.html	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/pig-0.12.0-cdh5.4.PMS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.5.4+0	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/search-1.0.0-cdh5.4.tar.gz	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/search-1.0.0-cdh5.4.html	https://us.apache.org/dist/hoodie/p/cdh5/cdh/5/search-1.0.0-cdh5.4.PMS.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.5.1+cdh5.5.4+138	https://us.apache.org/dist/hoodie/p/cdh5/cdh/	https://us.apache.org/dist/hoodie/p/cdh5/cdh/	https://us.apache.org/dist/hoodie/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ solr-4.10.3+cdh5.5.4.tar.gz	5/ solr-4.10.3+cdh5.5.4.html	5/ solr-4.10.3+cdh5.5.4.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.5.4+368	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.10.3+cdh5.5.4.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.10.3+cdh5.5.4.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.10.3+cdh5.5.4.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.5.0+cdh5.5.4+133	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-1.5.0+cdh5.5.4.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-1.5.0+cdh5.5.4.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-1.5.0+cdh5.5.4.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.5.4+41	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.6+cdh5.5.4.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.6+cdh5.5.4.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.6+cdh5.5.4.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.5.4+42	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.5+cdh5.5.4.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.5+cdh5.5.4.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.5+cdh5.5.4.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.5.4+23	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.5.4.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.5.4.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.5.4.html#changes
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.5.4+97	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.5.4.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.5.4.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.5.4.html#changes
		Tarball	Release notes	Changes

CDH 5.5.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.5.2+113	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.5.2+113.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.5.2+113.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.5.2+113.html#changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/arc1.76cdh5.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/arc1.76cdh5.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/arc1.76cdh5.2.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.5.2+78	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/crunch0.11.0cdh5.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/crunch0.11.0cdh5.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/crunch0.11.0cdh5.2.html#changes
		Tarball	Release notes	Changes
Apache DataFu (Incubating)	pig-df-datafu-1.10+cdh5.5.2+18	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/datafu1.10cdh5.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/datafu1.10cdh5.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/datafu1.10cdh5.2.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.6.0+cdh5.5.2+31	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume1.6.0cdh5.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume1.6.0cdh5.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume1.6.0cdh5.2.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.5.2+992	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop2.6.0cdh5.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop2.6.0cdh5.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop2.6.0cdh5.2.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.5.2+297	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.0.0cdh5.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.0.0cdh5.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.0.0cdh5.2.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.5.2+58	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.5cdh5.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.5cdh5.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.5cdh5.2.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.5.2+377	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive1.1.0cdh5.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive1.1.0cdh5.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive1.1.0cdh5.2.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Hue	hue-3.9.0+cdh5.5.2+362	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/hue390cdh52tar.gz	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/hue390cdh52releasenotes.html	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/hue390cdh52changes.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.3.0+cdh5.5.2+0	(none)	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/impala230cdh52releasenotes.html	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/impala230cdh52changes.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.5.2+120	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/kite100cdh52tar.gz	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/kite100cdh52releasenotes.html	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/kite100cdh52changes.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.5.2+0	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/llama100cdh52tar.gz	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/llama100cdh52releasenotes.html	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/llama100cdh52changes.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.5.2+27	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/mahout09cdh52tar.gz	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/mahout09cdh52releasenotes.html	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/mahout09cdh52changes.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.5.2+233	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/oozie410cdh52tar.gz	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/oozie410cdh52releasenotes.html	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/oozie410cdh52changes.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.5.2+175	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/parquet150cdh52tar.gz	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/parquet150cdh52releasenotes.html	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/parquet150cdh52changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.10.0+cdh5.5.2+13	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/parquet-format110cdh52tar.gz	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/parquet-format110cdh52releasenotes.html	https://s3.amazonaws.com/elasticdistribution/p/cdh5/cdh/5/parquet-format110cdh52changes.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.5.2+75	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh552tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh552releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh552changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.5.2+0	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh552tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh552releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh552changes.txt
		Tarball	Release notes	Changes
Apache Sentry (Incubating)	sentry-1.5.1+cdh5.5.2+120	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry151cdh552tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry151cdh552releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry151cdh552changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.5.2+340	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh552tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh552releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh552changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.5.0+cdh5.5.2+114	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark150cdh552tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark150cdh552releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark150cdh552changes.txt
		Tarball	Release notes	Changes
Spark-netlib	spark-netlib-master.2	(none)	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sparknetlib1cdh552releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sparknetlib1cdh552changes.txt
			Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.5.2+35	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop146cdh552tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop146cdh552releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop146cdh552changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.5.2+35	https://archive.apache.org/dist/parcels/p/cdh5/cdh/	https://archive.apache.org/dist/parcels/p/cdh5/cdh/	https://archive.apache.org/dist/parcels/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ spp-090-cdh5.5.2.tar.gz	5/ spp-090-cdh5.5.2.html	5/ spp-090-cdh5.5.2.html
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.5.2+18	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-090-cdh5.5.2.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-090-cdh5.5.2.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-090-cdh5.5.2.html
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.5.2+92	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-345-cdh5.5.2.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-345-cdh5.5.2.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-345-cdh5.5.2.html
		Tarball	Release notes	Changes

CDH 5.5.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.5.1+111	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/avro-176-cdh5.5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/avro-176-cdh5.5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/avro-176-cdh5.5.1.html
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.5.1+77	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/crunch-0110-cdh5.5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/crunch-0110-cdh5.5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/crunch-0110-cdh5.5.1.html
		Tarball	Release notes	Changes
DataFu	pig-of-datafu-1.10+cdh5.5.1+17	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/datafu-110-cdh5.5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/datafu-110-cdh5.5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/datafu-110-cdh5.5.1.html
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.6.0+cdh5.5.1+29	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/flume-ng-160-cdh5.5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/flume-ng-160-cdh5.5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/flume-ng-160-cdh5.5.1.html
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.5.1+924	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hadoop-260-cdh5.5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hadoop-260-cdh5.5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hadoop-260-cdh5.5.1.html
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/imp260cdh5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/imp260cdh5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/imp260cdh5.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.5.1+274	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.0.0cdh5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.0.0cdh5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.0.0cdh5.1.CHANGES.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.5.1+57	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.5cdh5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.5cdh5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.5cdh5.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.5.1+327	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive110cdh5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive110cdh5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive110cdh5.1.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.5.1+333	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue390cdh5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue390cdh5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue390cdh5.1.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.3.0+cdh5.5.1+0	(none)	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/imp230cdh5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/imp230cdh5.1.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.5.1+116	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite100cdh5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite100cdh5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite100cdh5.1.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.5.1+0	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama100cdh5.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama100cdh5.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama100cdh5.1.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Mahout	mahout-0.9+cdh5.5.1+26	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.5.1.tar.gz	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.5.1.html	https://archive.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.5.1Changes.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.5.1+223	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.5.1.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.5.1.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.5.1Changes.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.5.1+176	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.5.1.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.5.1.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.5.1Changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.5.1+12	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.5.1.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.5.1.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.5.1Changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.5.1+72	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.5.1.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.5.1.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.5.1Changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.5.1+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.5.1.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.5.1.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.5.1Changes.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.5.1+cdh5.5.1+106	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.5.1.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.5.1.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.5.1+cdh5.5.1Changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.5.1+325	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.5.1.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.5.1.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.5.1Changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Spark	spark-1.5.0+cdh5.5.1+94	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark150cdh551.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark150cdh551.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark150cdh551Changes.txt
		Tarball	Release notes	Changes
Spark Netlib	spark-netlib-master.3	(none)	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sparknetlib.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sparknetlibChanges.txt
			Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.5.1+29	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop146cdh551.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop146cdh551.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop146cdh551Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.5.1+33	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh551.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh551.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh551Changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.5.1+17	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr090cdh551.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr090cdh551.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr090cdh551Changes.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.5.1+91	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper345cdh551.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper345cdh551.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper345cdh551Changes.txt
		Tarball	Release notes	Changes

CDH 5.5.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.5.0+108	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/avro176cdh550.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/avro176cdh550.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/avro176cdh550Changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Crunch	crunch-0.11.0+cdh5.5.0+76	https://archive.apache.org/dist/hadoop/core/hadoop-0.11.0-cdh5.5.0.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-0.11.0-cdh5.5.0-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-0.11.0-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.5.0+16	https://archive.apache.org/dist/hadoop/pig/pig-1.10-cdh5.5.0.tar.gz	https://archive.apache.org/dist/hadoop/pig/pig-1.10-cdh5.5.0-release.html	https://archive.apache.org/dist/hadoop/pig/pig-1.10-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.6.0+cdh5.5.0+27	https://archive.apache.org/dist/hadoop/flume/flume-ng-1.6.0-cdh5.5.0.tar.gz	https://archive.apache.org/dist/hadoop/flume/flume-ng-1.6.0-cdh5.5.0-release.html	https://archive.apache.org/dist/hadoop/flume/flume-ng-1.6.0-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.5.0+921	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.5.0.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.5.0-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.5.0+269	https://archive.apache.org/dist/hadoop/hbase/hbase-1.0.0-cdh5.5.0.tar.gz	https://archive.apache.org/dist/hadoop/hbase/hbase-1.0.0-cdh5.5.0-release.html	https://archive.apache.org/dist/hadoop/hbase/hbase-1.0.0-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.5.0+56	https://archive.apache.org/dist/hadoop/hbase/hbase-solr-1.5-cdh5.5.0.tar.gz	https://archive.apache.org/dist/hadoop/hbase/hbase-solr-1.5-cdh5.5.0-release.html	https://archive.apache.org/dist/hadoop/hbase/hbase-solr-1.5-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.5.0+323	https://archive.apache.org/dist/hadoop/hive/hive-1.1.0-cdh5.5.0.tar.gz	https://archive.apache.org/dist/hadoop/hive/hive-1.1.0-cdh5.5.0-release.html	https://archive.apache.org/dist/hadoop/hive/hive-1.1.0-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Hue	hue-3.9.0+cdh5.5.0+332	https://archive.apache.org/dist/hadoop/hue/hue-3.9.0-cdh5.5.0.tar.gz	https://archive.apache.org/dist/hadoop/hue/hue-3.9.0-cdh5.5.0-release.html	https://archive.apache.org/dist/hadoop/hue/hue-3.9.0-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.3.0+cdh5.5.0+0	(none)	https://userguide.cloudera.com/p/cdh5/cdh/5/impala-230-cdh50.html	https://userguide.cloudera.com/p/cdh5/cdh/5/impala-230-cdh50.html#changes
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.5.0+115	https://userguide.cloudera.com/p/cdh5/cdh/5/kite-100-cdh50.tar.gz	https://userguide.cloudera.com/p/cdh5/cdh/5/kite-100-cdh50.html	https://userguide.cloudera.com/p/cdh5/cdh/5/kite-100-cdh50.html#changes
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.5.0+0	https://userguide.cloudera.com/p/cdh5/cdh/5/llama-100-cdh50.tar.gz	https://userguide.cloudera.com/p/cdh5/cdh/5/llama-100-cdh50.html	https://userguide.cloudera.com/p/cdh5/cdh/5/llama-100-cdh50.html#changes
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.5.0+25	https://userguide.cloudera.com/p/cdh5/cdh/5/mahout-09-cdh50.tar.gz	https://userguide.cloudera.com/p/cdh5/cdh/5/mahout-09-cdh50.html	https://userguide.cloudera.com/p/cdh5/cdh/5/mahout-09-cdh50.html#changes
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.5.0+222	https://userguide.cloudera.com/p/cdh5/cdh/5/oozie-410-cdh50.tar.gz	https://userguide.cloudera.com/p/cdh5/cdh/5/oozie-410-cdh50.html	https://userguide.cloudera.com/p/cdh5/cdh/5/oozie-410-cdh50.html#changes
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.5.0+173	https://userguide.cloudera.com/p/cdh5/cdh/5/parquet-150-cdh50.tar.gz	https://userguide.cloudera.com/p/cdh5/cdh/5/parquet-150-cdh50.html	https://userguide.cloudera.com/p/cdh5/cdh/5/parquet-150-cdh50.html#changes
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.5.0+11	https://userguide.cloudera.com/p/cdh5/cdh/5/parquet-format-210-cdh50.tar.gz	https://userguide.cloudera.com/p/cdh5/cdh/5/parquet-format-210-cdh50.html	https://userguide.cloudera.com/p/cdh5/cdh/5/parquet-format-210-cdh50.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.5.0+70	https://userguide.cloudera.com/p/cdh5/cdh/	https://userguide.cloudera.com/p/cdh5/cdh/	https://userguide.cloudera.com/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		https://archive.apache.org/dist/parcels/5/parcels-5.0.0-120-cdh5.5.0.tar.gz	https://archive.apache.org/dist/parcels/5/parcels-5.0.0-120-cdh5.5.0-releasenotes.html	https://archive.apache.org/dist/parcels/5/parcels-5.0.0-120-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.5.0+0	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search-1.0.0-cdh5.5.0.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search-1.0.0-cdh5.5.0-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search-1.0.0-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.5.1+cdh5.5.0+105	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.5.1-cdh5.5.0.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.5.1-cdh5.5.0-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.5.1-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.5.0+323	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.10.3-cdh5.5.0.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.10.3-cdh5.5.0-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.10.3-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.5.0+cdh5.5.0+92	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.5.0-cdh5.5.0.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.5.0-cdh5.5.0-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.5.0-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.6+cdh5.5.0+28	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.6-cdh5.5.0.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.6-cdh5.5.0-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.6-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.5.0+32	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.5-cdh5.5.0.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.5-cdh5.5.0-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.5-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.5.0+15	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0-cdh5.5.0.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0-cdh5.5.0-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0-cdh5.5.0-changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache ZooKeeper	zookeeper-3.4.5-cdh5.5.0+90	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5-cdh5.5.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5-cdh5.5.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5-cdh5.5.0.txt
		Tarball	Release notes	Changes

CDH 5.4.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.4.x release are listed in the following table. For the same type of information for other CDH releases, see [Packaging and Tarballs](#). To view the overall release notes for CDH 5.x.x, see the [Release Notes](#).



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.4.11 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.4.11+102	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.4.11.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.4.11.html	https://us.apache.org/dist/cdh5/cdh/5/avro-1.7.6+cdh5.4.11Changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.11+79	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.4.11.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.4.11.html	https://us.apache.org/dist/cdh5/cdh/5/crunch-0.11.0+cdh5.4.11Changes.txt
		Tarball	Release notes	Changes
Datafu	pig-udf-datafu-1.1.0+cdh5.4.11+27	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0+cdh5.4.11.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0+cdh5.4.11.html	https://us.apache.org/dist/cdh5/cdh/5/datafu-1.1.0+cdh5.4.11Changes.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.5.0+cdh5.4.11+146	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.5.0+cdh5.4.11.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.5.0+cdh5.4.11.html	https://us.apache.org/dist/cdh5/cdh/5/flume-ng-1.5.0+cdh5.4.11Changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.11+802	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.11.tar.gz	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.11.html	https://us.apache.org/dist/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.11Changes.txt
		Tarball	Release notes	Changes
Hadoop Mrv1	hadoop-0.20-mapreduce-2.6.0+cdh5.4.11+802	(none)	(none)	(none)

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Hbase	hbase-1.0.0+cdh5.4.11+245	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.11.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.11.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.11.CHANGES.txt
		Tarball	Release notes	Changes
Hbase-solr	hbase-solr-1.5+cdh5.4.11+64	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.4.11.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.4.11.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hbase-1.5+cdh5.4.11.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.11+350	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.11.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.11.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.11.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.4.11+1337	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.11.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.11.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.11.CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.2.0+cdh5.4.11+0	(none)	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.11.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.11.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.11+48	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.11.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.11.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.11.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.11+0	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.11.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.11.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.11.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.11+33	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.4.11.tar.gz	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.4.11.release.html	https://usapoc01.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.4.11.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.4.11+164	https://usapocdn.azureedge.net/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.11.tar.gz	https://usapocdn.azureedge.net/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.11-release.html	https://usapocdn.azureedge.net/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.11-CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.4.11+104	https://usapocdn.azureedge.net/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.11.tar.gz	https://usapocdn.azureedge.net/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.11-release.html	https://usapocdn.azureedge.net/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.11-CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.4.11+18	https://usapocdn.azureedge.net/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.4.11.tar.gz	https://usapocdn.azureedge.net/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.4.11-release.html	https://usapocdn.azureedge.net/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.4.11-CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.11+72	https://usapocdn.azureedge.net/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.11.tar.gz	https://usapocdn.azureedge.net/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.11-release.html	https://usapocdn.azureedge.net/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.11-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.11+0	https://usapocdn.azureedge.net/p/cdh5/cdh/5/search-1.0.0+cdh5.4.11.tar.gz	https://usapocdn.azureedge.net/p/cdh5/cdh/5/search-1.0.0+cdh5.4.11-release.html	https://usapocdn.azureedge.net/p/cdh5/cdh/5/search-1.0.0+cdh5.4.11-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry	sentry-1.4.0+cdh5.4.11+214	https://usapocdn.azureedge.net/p/cdh5/cdh/5/sentry-1.4.0+cdh5.4.11.tar.gz	https://usapocdn.azureedge.net/p/cdh5/cdh/5/sentry-1.4.0+cdh5.4.11-release.html	https://usapocdn.azureedge.net/p/cdh5/cdh/5/sentry-1.4.0+cdh5.4.11-CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.4.11+327	https://usapocdn.azureedge.net/p/cdh5/cdh/5/solr-4.10.3+cdh5.4.11.tar.gz	https://usapocdn.azureedge.net/p/cdh5/cdh/5/solr-4.10.3+cdh5.4.11-release.html	https://usapocdn.azureedge.net/p/cdh5/cdh/5/solr-4.10.3+cdh5.4.11-CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.3.0+cdh5.4.11+59	https://usapocdn.azureedge.net/p/cdh5/cdh/5/spark-1.3.0+cdh5.4.11.tar.gz	https://usapocdn.azureedge.net/p/cdh5/cdh/5/spark-1.3.0+cdh5.4.11-release.html	https://usapocdn.azureedge.net/p/cdh5/cdh/5/spark-1.3.0+cdh5.4.11-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Sqoop	sqoop-1.4.5+cdh5.4.11+125	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.11.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.11.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.11Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.4.11+46	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.11.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.11.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.11Changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.4.11+24	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.4.11.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.4.11.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.4.11Changes.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.4.11+100	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.4.11.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.4.11.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zoo-3.4.5+cdh5.4.11Changes.txt
		Tarball	Release notes	Changes

CDH 5.4.10 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.4.10+101	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.4.10.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.4.10.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/av-1.7.6+cdh5.4.10Changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.10+78	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.4.10.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.4.10.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/crunch-0.11.0+cdh5.4.10Changes.txt
		Tarball	Release notes	Changes
DataFu	pig-dbfu-1.10+cdh5.4.10+26	https://archive.apache.org/dist/parcels/cdh5/cdh/5/dbfu-1.10+cdh5.4.10.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/dbfu-1.10+cdh5.4.10.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/dbfu-1.10+cdh5.4.10Changes.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.5.0+cdh5.4.10+143	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.5.0+cdh5.4.10.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.5.0+cdh5.4.10.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/flume-ng-1.5.0+cdh5.4.10Changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.4.10.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.4.10-release.html	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.4.10-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.10+746	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.4.10.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.4.10-release.html	https://us.apache.org/dist/hadoop/core/hadoop-2.6.0+cdh5.4.10-changes.txt
		Tarball	Release notes	Changes
HBase	hbase-1.0.0+cdh5.4.10+226	https://us.apache.org/dist/hadoop/hbase/hbase-1.0.0+cdh5.4.10.tar.gz	https://us.apache.org/dist/hadoop/hbase/hbase-1.0.0+cdh5.4.10-release.html	https://us.apache.org/dist/hadoop/hbase/hbase-1.0.0+cdh5.4.10-changes.txt
		Tarball	Release notes	Changes
HBase-solr	hbase-solr-1.5+cdh5.4.10+62	https://us.apache.org/dist/hadoop/hbase/hbase-solr-1.5+cdh5.4.10.tar.gz	https://us.apache.org/dist/hadoop/hbase/hbase-solr-1.5+cdh5.4.10-release.html	https://us.apache.org/dist/hadoop/hbase/hbase-solr-1.5+cdh5.4.10-changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.10+315	https://us.apache.org/dist/hadoop/hive/hive-1.1.0+cdh5.4.10.tar.gz	https://us.apache.org/dist/hadoop/hive/hive-1.1.0+cdh5.4.10-release.html	https://us.apache.org/dist/hadoop/hive/hive-1.1.0+cdh5.4.10-changes.txt
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.4.10+1330	https://us.apache.org/dist/hadoop/hue/hue-3.7.0+cdh5.4.10.tar.gz	https://us.apache.org/dist/hadoop/hue/hue-3.7.0+cdh5.4.10-release.html	https://us.apache.org/dist/hadoop/hue/hue-3.7.0+cdh5.4.10-changes.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.2.0+cdh5.4.10+0	(none)	https://us.apache.org/dist/hadoop/impala/impala-2.2.0+cdh5.4.10-release.html	https://us.apache.org/dist/hadoop/impala/impala-2.2.0+cdh5.4.10-changes.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.10+47	https://us.apache.org/dist/hadoop/kite/kite-1.0.0+cdh5.4.10.tar.gz	https://us.apache.org/dist/hadoop/kite/kite-1.0.0+cdh5.4.10-release.html	https://us.apache.org/dist/hadoop/kite/kite-1.0.0+cdh5.4.10-changes.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.10+0	https://userguide.hortonworks.com/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.10.tar.gz	https://userguide.hortonworks.com/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.10.release.html	https://userguide.hortonworks.com/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.10.changes.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.10+32	https://userguide.hortonworks.com/p/cdh5/cdh/5/mahout-0.9+cdh5.4.10.tar.gz	https://userguide.hortonworks.com/p/cdh5/cdh/5/mahout-0.9+cdh5.4.10.release.html	https://userguide.hortonworks.com/p/cdh5/cdh/5/mahout-0.9+cdh5.4.10.changes.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.4.10+159	https://userguide.hortonworks.com/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.10.tar.gz	https://userguide.hortonworks.com/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.10.release.html	https://userguide.hortonworks.com/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.10.changes.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.4.10+103	https://userguide.hortonworks.com/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.10.tar.gz	https://userguide.hortonworks.com/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.10.release.html	https://userguide.hortonworks.com/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.10.changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.10.0+cdh5.4.10+17	https://userguide.hortonworks.com/p/cdh5/cdh/5/parquet-format-2.10.0+cdh5.4.10.tar.gz	https://userguide.hortonworks.com/p/cdh5/cdh/5/parquet-format-2.10.0+cdh5.4.10.release.html	https://userguide.hortonworks.com/p/cdh5/cdh/5/parquet-format-2.10.0+cdh5.4.10.changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.10+71	https://userguide.hortonworks.com/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.10.tar.gz	https://userguide.hortonworks.com/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.10.release.html	https://userguide.hortonworks.com/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.10.changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.10+0	https://userguide.hortonworks.com/p/cdh5/cdh/5/search-1.0.0+cdh5.4.10.tar.gz	https://userguide.hortonworks.com/p/cdh5/cdh/5/search-1.0.0+cdh5.4.10.release.html	https://userguide.hortonworks.com/p/cdh5/cdh/5/search-1.0.0+cdh5.4.10.changes.txt
		Tarball	Release notes	Changes
Apache Sentry (Incubating)	sentry-1.4.0+cdh5.4.10+205	https://userguide.hortonworks.com/p/cdh5/cdh/	https://userguide.hortonworks.com/p/cdh5/cdh/	https://userguide.hortonworks.com/p/cdh5/cdh/

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		p/cdh5/cdh/5/arc1.76-cdh5.4.9.tar.gz	p/cdh5/cdh/5/arc1.76-cdh5.4.9.html	p/cdh5/cdh/5/arc1.76-cdh5.4.9.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.9+77	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/crunch0.11.0-cdh5.4.9.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/crunch0.11.0-cdh5.4.9.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/crunch0.11.0-cdh5.4.9.html#changes
		Tarball	Release notes	Changes
DataFu	pig-udf-datafu-1.10+cdh5.4.9+25	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/datafu1.10-cdh5.4.9.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/datafu1.10-cdh5.4.9.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/datafu1.10-cdh5.4.9.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.4.9+139	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume1.5.0-cdh5.4.9.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume1.5.0-cdh5.4.9.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume1.5.0-cdh5.4.9.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.9+693	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop2.6.0-cdh5.4.9.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop2.6.0-cdh5.4.9.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop2.6.0-cdh5.4.9.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.4.9+212	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.0.0-cdh5.4.9.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.0.0-cdh5.4.9.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.0.0-cdh5.4.9.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.4.9+61	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.5-cdh5.4.9.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.5-cdh5.4.9.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.5-cdh5.4.9.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.9+294	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive1.1.0-cdh5.4.9.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive1.1.0-cdh5.4.9.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive1.1.0-cdh5.4.9.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Hue	hue-3.7.0+cdh5.4.9+1314	https://s3.amazonaws.com/cdh5/cdh/5/hue370cdh549.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue370cdh549.html	https://s3.amazonaws.com/cdh5/cdh/5/hue370cdh549.PMS.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.2.0+cdh5.4.9+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/imp220cdh549.html	https://s3.amazonaws.com/cdh5/cdh/5/imp220cdh549.PMS.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.9+43	https://s3.amazonaws.com/cdh5/cdh/5/kite100cdh549.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite100cdh549.html	https://s3.amazonaws.com/cdh5/cdh/5/kite100cdh549.PMS.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.9+0	https://s3.amazonaws.com/cdh5/cdh/5/llama100cdh549.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama100cdh549.html	https://s3.amazonaws.com/cdh5/cdh/5/llama100cdh549.PMS.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.9+31	https://s3.amazonaws.com/cdh5/cdh/5/mahout09cdh549.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout09cdh549.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout09cdh549.PMS.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.4.9+153	https://s3.amazonaws.com/cdh5/cdh/5/oozie410cdh549.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie410cdh549.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie410cdh549.PMS.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.4.9+102	https://s3.amazonaws.com/cdh5/cdh/5/parquet150cdh549.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet150cdh549.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet150cdh549.PMS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.10+cdh5.4.9+16	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format210cdh549.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format210cdh549.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format210cdh549.PMS.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.9+70	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh549.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh549.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh549CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.9+0	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh549.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh549.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh549CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.4.9+191	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry140cdh549.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry140cdh549.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry140cdh549CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.4.9+299	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh549.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh549.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh549CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.3.0+cdh5.4.9+54	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark130cdh549.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark130cdh549.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark130cdh549CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.4.9+120	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop145cdh549.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop145cdh549.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop145cdh549CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.4.9+41	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh549.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh549.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh549CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.4.9+22	https://archive.apache.org/dist/parcels/p/cdh5/cdh/	https://archive.apache.org/dist/parcels/p/cdh5/cdh/	https://archive.apache.org/dist/parcels/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ zip-090cdh549.tar.gz	5/ zip-090cdh549.html	5/ zip-090cdh549.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.4.9+97	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ zip-345cdh549.tar.gz	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ zip-345cdh549.html	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ zip-345cdh549.txt
		Tarball	Release notes	Changes

CDH 5.4.8 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.4.8+97	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ avro-1.7.6cdh548.tar.gz	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ avro-1.7.6cdh548.html	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ avro-1.7.6cdh548.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.8+76	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ crunch0110cdh548.tar.gz	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ crunch0110cdh548.html	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ crunch0110cdh548.txt
		Tarball	Release notes	Changes
DataFu	pig-df-datafu-1.10+cdh5.4.8+24	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ datafu-1.10cdh548.tar.gz	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ datafu-1.10cdh548.html	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ datafu-1.10cdh548.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.4.8+137	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ flume-ng-1.5.0cdh548.tar.gz	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ flume-ng-1.5.0cdh548.html	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ flume-ng-1.5.0cdh548.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.8+669	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ hadoop-2.6.0cdh548.tar.gz	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ hadoop-2.6.0cdh548.html	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ hadoop-2.6.0cdh548.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.4.8+195	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ hbase-1.0.0cdh548.tar.gz	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ hbase-1.0.0cdh548.html	https:// www.apache.org/dist/cdh/ p/cdh5/cdh/ 5/ hbase-1.0.0cdh548.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.4.8+60	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.8+275	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.4.8+1303	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.2.0+cdh5.4.8+0	(none)	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.8.html#changes
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.8+41	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.8+0	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.8+30	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/mahout-0.9+cdh5.4.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/mahout-0.9+cdh5.4.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/mahout-0.9+cdh5.4.8.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.4.8+152	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.8.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.4.8+99	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.8.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.10+cdh5.4.8+15	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.10+cdh5.4.8.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.10+cdh5.4.8.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.10+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.8+67	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.8.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.8+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.4.8.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.4.8+183	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.4.8.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.4.8+287	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.4.8.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.4.8.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.4.8.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.3.0+cdh5.4.8+51	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.3.0+cdh5.4.8.tar.gz	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.3.0+cdh5.4.8.html	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.3.0+cdh5.4.8.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.4.8+119	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.8.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.8.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.8.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.4.8+40	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.8.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.8.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.8.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.4.8+20	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/whirr-0.9.0+cdh5.4.8.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/whirr-0.9.0+cdh5.4.8.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/whirr-0.9.0+cdh5.4.8.CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.4.8+96	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.4.8.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.4.8.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.4.8.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.4.7 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.4.7+96	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/avro-1.7.6+cdh5.4.7.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/avro-1.7.6+cdh5.4.7.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/avro-1.7.6+cdh5.4.7.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.7+75	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/crunch-0.11.0+cdh5.4.7.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/crunch-0.11.0+cdh5.4.7.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/crunch-0.11.0+cdh5.4.7.CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.4.7+24	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/datafu-1.10+cdh5.4.7.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/datafu-1.10+cdh5.4.7.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/datafu-1.10+cdh5.4.7.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Flume	flume-ng-1.5.0+cdh5.4.7+136	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-1.5.0+cdh5.4.7.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-1.5.0+cdh5.4.7.release.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-1.5.0+cdh5.4.7.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.7+642	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.7.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.7.release.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.7.CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.4.7+183	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.7.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.7.release.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.7.CHANGES.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.4.7+60	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.7.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.7.release.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.7.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.7+233	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.7.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.7.release.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.7.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.4.7+1298	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.7.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.7.release.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.7.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.2.0+cdh5.4.7+0	(none)	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.7.release.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.7.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.7+39	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.7.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.7.release.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.7.CHANGES.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.7+0	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.7.tar.gz	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.7.release.html	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.7.PMS.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.7+27	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/mahout-0.9-cdh5.4.7.tar.gz	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/mahout-0.9-cdh5.4.7.release.html	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/mahout-0.9-cdh5.4.7.PMS.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.4.7+151	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/oozie-4.1.0-cdh5.4.7.tar.gz	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/oozie-4.1.0-cdh5.4.7.release.html	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/oozie-4.1.0-cdh5.4.7.PMS.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.4.7+99	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/parquet-1.5.0-cdh5.4.7.tar.gz	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/parquet-1.5.0-cdh5.4.7.release.html	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/parquet-1.5.0-cdh5.4.7.PMS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.10+cdh5.4.7+15	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/parquet-format-2.10-cdh5.4.7.tar.gz	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/parquet-format-2.10-cdh5.4.7.release.html	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/parquet-format-2.10-cdh5.4.7.PMS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.7+67	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/pig-0.12.0-cdh5.4.7.tar.gz	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/pig-0.12.0-cdh5.4.7.release.html	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/pig-0.12.0-cdh5.4.7.PMS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.7+0	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/search-1.0.0-cdh5.4.7.tar.gz	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/search-1.0.0-cdh5.4.7.release.html	https://us.apache.org/dist/traffic/p/cdh5/cdh/5/search-1.0.0-cdh5.4.7.PMS.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.4.7+179	https://us.apache.org/dist/traffic/p/cdh5/cdh/	https://us.apache.org/dist/traffic/p/cdh5/cdh/	https://us.apache.org/dist/traffic/p/cdh5/cdh/

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/arc1.76cdh5.45.tar.gz	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/arc1.76cdh5.45.html	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/arc1.76cdh5.45.PDF.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.5+58	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/crunch0.11.0cdh5.45.tar.gz	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/crunch0.11.0cdh5.45.html	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/crunch0.11.0cdh5.45.PDF.txt
		Tarball	Release notes	Changes
DataFu	pig-udf-datafu-1.10+cdh5.4.5+22	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/datafu1.10cdh5.45.tar.gz	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/datafu1.10cdh5.45.html	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/datafu1.10cdh5.45.PDF.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.4.5+132	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/flume-ng1.5.0cdh5.45.tar.gz	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/flume-ng1.5.0cdh5.45.html	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/flume-ng1.5.0cdh5.45.PDF.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.5+626	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hadoop2.6.0cdh5.45.tar.gz	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hadoop2.6.0cdh5.45.html	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hadoop2.6.0cdh5.45.PDF.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.4.5+175	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hbase1.0.0cdh5.45.tar.gz	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hbase1.0.0cdh5.45.html	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hbase1.0.0cdh5.45.PDF.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.4.5+57	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hbase-solr1.5cdh5.45.tar.gz	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hbase-solr1.5cdh5.45.html	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hbase-solr1.5cdh5.45.PDF.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.5+199	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hive1.1.0cdh5.45.tar.gz	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hive1.1.0cdh5.45.html	https://s3.amazonaws.com/cdh-downloads-us-east-1/p/cdh5/cdh/5/hive1.1.0cdh5.45.PDF.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Hue	hue-3.7.0+cdh5.4.5+1272	https://s3.amazonaws.com/cdh5/cdh/5/hue370cdh545.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue370cdh545.html	https://s3.amazonaws.com/cdh5/cdh/5/hue370cdh545CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.2.0+cdh5.4.5+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/imp220cdh545.html	https://s3.amazonaws.com/cdh5/cdh/5/imp220cdh545CHANGES.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.5+34	https://s3.amazonaws.com/cdh5/cdh/5/kite100cdh545.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite100cdh545.html	https://s3.amazonaws.com/cdh5/cdh/5/kite100cdh545CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.5+0	https://s3.amazonaws.com/cdh5/cdh/5/llama100cdh545.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama100cdh545.html	https://s3.amazonaws.com/cdh5/cdh/5/llama100cdh545CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.5+25	https://s3.amazonaws.com/cdh5/cdh/5/mahout09cdh545.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout09cdh545.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout09cdh545CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.4.5+149	https://s3.amazonaws.com/cdh5/cdh/5/oozie410cdh545.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie410cdh545.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie410cdh545CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.4.5+97	https://s3.amazonaws.com/cdh5/cdh/5/parquet150cdh545.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet150cdh545.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet150cdh545CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.10+cdh5.4.5+13	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format210cdh545.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format210cdh545.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format210cdh545CHANGES.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.5+65	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh545.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh545.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh545CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.5+0	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh545.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh545.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh545CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.4.5+175	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry140cdh545.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry140cdh545.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry140cdh545CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.4.5+263	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh545.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh545.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh545CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.3.0+cdh5.4.5+47	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark130cdh545.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark130cdh545.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark130cdh545CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.4.5+114	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop145cdh545.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop145cdh545.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop145cdh545CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.4.5+38	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh545.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh545.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh545CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.4.5+18	https://archive.apache.org/dist/parcels/p/cdh5/cdh/	https://archive.apache.org/dist/parcels/p/cdh5/cdh/	https://archive.apache.org/dist/parcels/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ hadoop-core-0.90.0-cdh5.4.5.tar.gz	5/ hadoop-core-0.90.0-cdh5.4.5.html	5/ hadoop-core-0.90.0-cdh5.4.5.html
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.4.5+92	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ zookeeper-3.4.5-cdh5.4.5.tar.gz	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ zookeeper-3.4.5-cdh5.4.5.html	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ zookeeper-3.4.5-cdh5.4.5.html
		Tarball	Release notes	Changes

CDH 5.4.4 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.4.4+91	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ avro-1.7.6-cdh5.4.4.tar.gz	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ avro-1.7.6-cdh5.4.4.html	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ avro-1.7.6-cdh5.4.4.html
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.4+49	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ crunch-0.11.0-cdh5.4.4.tar.gz	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ crunch-0.11.0-cdh5.4.4.html	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ crunch-0.11.0-cdh5.4.4.html
		Tarball	Release notes	Changes
DataFu	pig-df-datafu-1.10+cdh5.4.4+19	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ datafu-1.10-cdh5.4.4.tar.gz	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ datafu-1.10-cdh5.4.4.html	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ datafu-1.10-cdh5.4.4.html
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.4.4+125	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ flume-ng-1.5.0-cdh5.4.4.tar.gz	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ flume-ng-1.5.0-cdh5.4.4.html	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ flume-ng-1.5.0-cdh5.4.4.html
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.4+597	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ hadoop-2.6.0-cdh5.4.4.tar.gz	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ hadoop-2.6.0-cdh5.4.4.html	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ hadoop-2.6.0-cdh5.4.4.html
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.4.4+160	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ hbase-1.0.0-cdh5.4.4.tar.gz	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ hbase-1.0.0-cdh5.4.4.html	https:// www.apache.org/dist/hadoop/ p/cdh5/cdh/ 5/ hbase-1.0.0-cdh5.4.4.html
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hbase-1.0.0-cdh5.4.4.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hbase-1.0.0-cdh5.4.4.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hbase-1.0.0-cdh5.4.4.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.4.4+53	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.4.4.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.4.4.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.4.4.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.4+152	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hive-1.1.0-cdh5.4.4.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hive-1.1.0-cdh5.4.4.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hive-1.1.0-cdh5.4.4.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.4.4+1236	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hue-3.7.0-cdh5.4.4.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hue-3.7.0-cdh5.4.4.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/hue-3.7.0-cdh5.4.4.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.2.0+cdh5.4.4+0	(none)	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/impala-2.2.0-cdh5.4.4.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/impala-2.2.0-cdh5.4.4.html#changes
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.4+25	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/kite-1.0.0-cdh5.4.4.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/kite-1.0.0-cdh5.4.4.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/kite-1.0.0-cdh5.4.4.html#changes
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.4+0	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.4.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.4.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.4.html#changes
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.4+24	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/mahout-0.9-cdh5.4.4.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/mahout-0.9-cdh5.4.4.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/mahout-0.9-cdh5.4.4.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.4.4+145	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/oozie-410-cdh544.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/oozie-410-cdh544.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/oozie-410-cdh544Changes.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.4.4+96	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-150-cdh544.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-150-cdh544.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-150-cdh544Changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-210+cdh5.4.4+12	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-format-210-cdh544.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-format-210-cdh544.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-format-210-cdh544Changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.4+59	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/pig-0120-cdh544.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/pig-0120-cdh544.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/pig-0120-cdh544Changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.4+0	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/search-100-cdh544.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/search-100-cdh544.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/search-100-cdh544Changes.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.4.4+169	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/sentry-140-cdh544.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/sentry-140-cdh544.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/sentry-140-cdh544Changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.4.4+257	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/solr-4103-cdh544.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/solr-4103-cdh544.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/solr-4103-cdh544Changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.3.0+cdh5.4.4+41	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/spark-130-cdh544.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/spark-130-cdh544.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/spark-130-cdh544Changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.4.4+101	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.4.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.4.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.4.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.4.4+36	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.4.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.4.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.4.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.4.4+17	https://usapoc01r1zrh0em/p/cdh5/cdh/5/whirr-0.9.0+cdh5.4.4.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/whirr-0.9.0+cdh5.4.4.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/whirr-0.9.0+cdh5.4.4.CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.4.4+91	https://usapoc01r1zrh0em/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.4.4.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.4.4.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.4.4.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.4.3 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.4.3+90	https://usapoc01r1zrh0em/p/cdh5/cdh/5/avro-1.7.6+cdh5.4.3.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/avro-1.7.6+cdh5.4.3.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/avro-1.7.6+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.3+48	https://usapoc01r1zrh0em/p/cdh5/cdh/5/crunch-0.11.0+cdh5.4.3.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/crunch-0.11.0+cdh5.4.3.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/crunch-0.11.0+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.4.3+18	https://usapoc01r1zrh0em/p/cdh5/cdh/5/datafu-1.10+cdh5.4.3.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/datafu-1.10+cdh5.4.3.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/datafu-1.10+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Flume	flume-ng-1.5.0+cdh5.4.3+124	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/flume-1.5.0+cdh5.4.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/flume-1.5.0+cdh5.4.3-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/flume-1.5.0+cdh5.4.3-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.3+596	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.3-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.3-changes.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.4.3+159	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.3-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.3-changes.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.4.3+52	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.3-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.3-changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.3+151	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.3-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.3-changes.txt
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.4.3+1234	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.3-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.3-changes.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.2.0+cdh5.4.3+0	(none)	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.3-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.3-changes.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.3+24	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.3-release.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.6.0/cdh5/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.3-changes.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.3+0	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.3.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.3.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.3+23	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.4.3.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.4.3.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout-0.9+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.4.3+144	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.3.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.3.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.4.3+95	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.3.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.3.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.10+cdh5.4.3+11	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-2.10+cdh5.4.3.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-2.10+cdh5.4.3.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format-2.10+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.3+58	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.3.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.3.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.3+0	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/search-1.0.0+cdh5.4.3.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/search-1.0.0+cdh5.4.3.release.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/search-1.0.0+cdh5.4.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.4.3+168	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ solr-4.10.3+cdh5.4.3.tar.gz	5/ solr-4.10.3+cdh5.4.3.html	5/ solr-4.10.3+cdh5.4.3.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.4.3+256	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.10.3+cdh5.4.3.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.10.3+cdh5.4.3.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.10.3+cdh5.4.3.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.3.0+cdh5.4.3+40	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-1.3.0+cdh5.4.3.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-1.3.0+cdh5.4.3.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-1.3.0+cdh5.4.3.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.4.3+100	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.4.3.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.4.3.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.4.3.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.4.3+35	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.5+cdh5.4.3.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.5+cdh5.4.3.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.5+cdh5.4.3.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.4.3+16	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.4.3.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.4.3.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.4.3.html#changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.4.3+90	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.4.3.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.4.3.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.4.3.html#changes
		Tarball	Release notes	Changes

CDH 5.4.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.4.2+88	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.4.2+88.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.4.2+88.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.4.2+88.html#changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		p/cdh5/cdh/5/arc1.76-cdh5.4.2.tar.gz	p/cdh5/cdh/5/arc1.76-cdh5.4.2.html	p/cdh5/cdh/5/arc1.76-cdh5.4.2.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.2+46	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/crunch0.11.0-cdh5.4.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/crunch0.11.0-cdh5.4.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/crunch0.11.0-cdh5.4.2.html#changes
		Tarball	Release notes	Changes
DataFu	pig-udf-datafu-1.10+cdh5.4.2+16	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/datafu1.10-cdh5.4.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/datafu1.10-cdh5.4.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/datafu1.10-cdh5.4.2.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.4.2+122	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume1.5.0-cdh5.4.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume1.5.0-cdh5.4.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume1.5.0-cdh5.4.2.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.2+572	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop2.6.0-cdh5.4.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop2.6.0-cdh5.4.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop2.6.0-cdh5.4.2.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.4.2+143	(none)	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.0.0-cdh5.4.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.0.0-cdh5.4.2.html#changes
			Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.4.2+50	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.5-cdh5.4.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.5-cdh5.4.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase1.5-cdh5.4.2.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.2+141	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive1.1.0-cdh5.4.2.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive1.1.0-cdh5.4.2.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive1.1.0-cdh5.4.2.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Hue	hue-3.7.0+cdh5.4.2+1220	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hue370cdh542.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hue370cdh542.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/hue370cdh542.PMS.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.2.0+cdh5.4.2+0	(none)	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/imp220cdh542.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/imp220cdh542.PMS.txt
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.2+22	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kite100cdh542.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kite100cdh542.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/kite100cdh542.PMS.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.2+0	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama100cdh542.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama100cdh542.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/llama100cdh542.PMS.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.2+21	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout09cdh542.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout09cdh542.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/mahout09cdh542.PMS.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.4.2+142	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie410cdh542.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie410cdh542.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/oozie410cdh542.PMS.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.4.2+93	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet150cdh542.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet150cdh542.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet150cdh542.PMS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.10+cdh5.4.2+9	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format210cdh542.tar.gz	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format210cdh542.html	https://usapocdl.blob.core.windows.net/p/cdh5/cdh/5/parquet-format210cdh542.PMS.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.2+56	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh542.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh542.release.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/pig0120cdh542.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.2+0	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh542.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh542.release.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/search100cdh542.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.4.2+165	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry140cdh542.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry140cdh542.release.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry140cdh542.CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.4.2+251	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh542.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh542.release.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr4103cdh542.CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.3.0+cdh5.4.2+38	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark130cdh542.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark130cdh542.release.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark130cdh542.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.4.2+98	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop145cdh542.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop145cdh542.release.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop145cdh542.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.4.2+32	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh542.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh542.release.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop21995cdh542.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.4.2+14	https://archive.apache.org/dist/parcels/p/cdh5/cdh/	https://archive.apache.org/dist/parcels/p/cdh5/cdh/	https://archive.apache.org/dist/parcels/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ http://www.apache.org/dist/hadoop/core/hadoop-core-0.90.0-cdh5.4.2.tar.gz	5/ http://www.apache.org/dist/hadoop/core/hadoop-core-0.90.0-cdh5.4.2-release.html	5/ http://www.apache.org/dist/hadoop/core/hadoop-core-0.90.0-cdh5.4.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.4.2+88	https://www.apache.org/dist/hadoop/zookeeper/zookeeper-3.4.5-cdh5.4.2.tar.gz	https://www.apache.org/dist/hadoop/zookeeper/zookeeper-3.4.5-cdh5.4.2-release.html	https://www.apache.org/dist/hadoop/zookeeper/zookeeper-3.4.5-cdh5.4.2-CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.4.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.4.1+88	https://www.apache.org/dist/hadoop/avro/avro-1.7.6-cdh5.4.1.tar.gz	https://www.apache.org/dist/hadoop/avro/avro-1.7.6-cdh5.4.1-release.html	https://www.apache.org/dist/hadoop/avro/avro-1.7.6-cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.1+46	https://www.apache.org/dist/hadoop/crunch/crunch-0.11.0-cdh5.4.1.tar.gz	https://www.apache.org/dist/hadoop/crunch/crunch-0.11.0-cdh5.4.1-release.html	https://www.apache.org/dist/hadoop/crunch/crunch-0.11.0-cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pig-df-datafu-1.10+cdh5.4.1+16	https://www.apache.org/dist/hadoop/datafu/datafu-1.10-cdh5.4.1.tar.gz	https://www.apache.org/dist/hadoop/datafu/datafu-1.10-cdh5.4.1-release.html	https://www.apache.org/dist/hadoop/datafu/datafu-1.10-cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.4.1+121	https://www.apache.org/dist/hadoop/flume/flume-ng-1.5.0-cdh5.4.1.tar.gz	https://www.apache.org/dist/hadoop/flume/flume-ng-1.5.0-cdh5.4.1-release.html	https://www.apache.org/dist/hadoop/flume/flume-ng-1.5.0-cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.1+565	https://www.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.4.1.tar.gz	https://www.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.4.1-release.html	https://www.apache.org/dist/hadoop/core/hadoop-2.6.0-cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.4.1+141	https://www.apache.org/dist/hadoop/hbase/hbase-1.0.0-cdh5.4.1.tar.gz	https://www.apache.org/dist/hadoop/hbase/hbase-1.0.0-cdh5.4.1-release.html	https://www.apache.org/dist/hadoop/hbase/hbase-1.0.0-cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.1.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.4.1+50	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.1.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.1+126	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.1.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.4.1+1173	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.1.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.2.0+cdh5.4.1+0	(none)	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.1.html#changes
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.1+17	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.1.html#changes
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.1+0	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/llama-1.0.0+cdh5.4.1.html#changes
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.1+21	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/mahout-0.9+cdh5.4.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/mahout-0.9+cdh5.4.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/mahout-0.9+cdh5.4.1.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Oozie	oozie-4.1.0+cdh5.4.1+142	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.1-release.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0+cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.4.1+93	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.1-release.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.4.1+9	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.4.1-release.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.1+56	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.1-release.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.1+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.4.1-release.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.4.1+163	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.4.1-release.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.4.1+251	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.4.1.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.4.1-release.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.10.3+cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.3.0+cdh5.4.1+35	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.3.0+cdh5.4.1.tar.gz	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.3.0+cdh5.4.1-release.html	https://archive.apache.org/dist/spark/p/cdh5/cdh/5/spark-1.3.0+cdh5.4.1-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.4.1+98	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.1.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.1.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.4.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.4.1+28	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.1.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.1.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/sqoop2-1.99.5+cdh5.4.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.4.1+14	https://usapoc01r1zrh0em/p/cdh5/cdh/5/whirr-0.9.0+cdh5.4.1.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/whirr-0.9.0+cdh5.4.1.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/whirr-0.9.0+cdh5.4.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.4.1+87	https://usapoc01r1zrh0em/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.4.1.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.4.1.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.4.1.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.4.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.4.0+87	https://usapoc01r1zrh0em/p/cdh5/cdh/5/avro-1.7.6+cdh5.4.0.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/avro-1.7.6+cdh5.4.0.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/avro-1.7.6+cdh5.4.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.4.0+45	https://usapoc01r1zrh0em/p/cdh5/cdh/5/crunch-0.11.0+cdh5.4.0.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/crunch-0.11.0+cdh5.4.0.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/crunch-0.11.0+cdh5.4.0.CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.4.0+15	https://usapoc01r1zrh0em/p/cdh5/cdh/5/datafu-1.10+cdh5.4.0.tar.gz	https://usapoc01r1zrh0em/p/cdh5/cdh/5/datafu-1.10+cdh5.4.0.html	https://usapoc01r1zrh0em/p/cdh5/cdh/5/datafu-1.10+cdh5.4.0.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Flume	flume-ng-1.5.0+cdh5.4.0+120	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.4.0.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-1.5.0+cdh5.4.0.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-1.5.0+cdh5.4.0.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.6.0+cdh5.4.0+521	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.0.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.0.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.6.0+cdh5.4.0.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-1.0.0+cdh5.4.0+130	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.0.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.0.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-1.0.0+cdh5.4.0.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.4.0+49	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.0.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.0.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.4.0.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-1.1.0+cdh5.4.0+103	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.0.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.0.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-1.1.0+cdh5.4.0.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.4.0+1145	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.0.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.0.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hue-3.7.0+cdh5.4.0.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.2.0+cdh5.4.0+0	(none)	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.0.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/impala-2.2.0+cdh5.4.0.html#changes
			Release notes	Changes
Kite SDK	kite-1.0.0+cdh5.4.0+15	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.0.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.0.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/kite-1.0.0+cdh5.4.0.html#changes
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.4.0+0	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.0.tar.gz	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.0.html	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/llama-1.0.0-cdh5.4.0.PMS.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.4.0+20	https://us.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9-cdh5.4.0.tar.gz	https://us.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9-cdh5.4.0.html	https://us.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9-cdh5.4.0.PMS.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.1.0+cdh5.4.0+138	https://us.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.4.0.tar.gz	https://us.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.4.0.html	https://us.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.1.0-cdh5.4.0.PMS.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.4.0+92	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0-cdh5.4.0.tar.gz	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0-cdh5.4.0.html	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0-cdh5.4.0.PMS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.4.0+8	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.4.0.tar.gz	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.4.0.html	https://us.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-2.1.0-cdh5.4.0.PMS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.4.0+55	https://us.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0-cdh5.4.0.tar.gz	https://us.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0-cdh5.4.0.html	https://us.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0-cdh5.4.0.PMS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.4.0+0	https://us.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0-cdh5.4.0.tar.gz	https://us.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0-cdh5.4.0.html	https://us.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0-cdh5.4.0.PMS.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.4.0+155	https://us.apache.org/dist/sentry/p/cdh5/cdh/	https://us.apache.org/dist/sentry/p/cdh5/cdh/	https://us.apache.org/dist/sentry/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ parquet-1.4.0-cdh5.4.0.tar.gz	5/ parquet-1.4.0-cdh5.4.0-release.html	5/ parquet-1.4.0-cdh5.4.0-changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.10.3+cdh5.4.0+240	https://archive.apache.org/dist/solr/solr-4.10.3-cdh5.4.0.tar.gz	https://archive.apache.org/dist/solr/solr-4.10.3-cdh5.4.0-release.html	https://archive.apache.org/dist/solr/solr-4.10.3-cdh5.4.0-changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.3.0+cdh5.4.0+24	https://archive.apache.org/dist/spark/spark-1.3.0-cdh5.4.0.tar.gz	https://archive.apache.org/dist/spark/spark-1.3.0-cdh5.4.0-release.html	https://archive.apache.org/dist/spark/spark-1.3.0-cdh5.4.0-changes.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.4.0+91	https://archive.apache.org/dist/sqoop/sqoop-1.4.5-cdh5.4.0.tar.gz	https://archive.apache.org/dist/sqoop/sqoop-1.4.5-cdh5.4.0-release.html	https://archive.apache.org/dist/sqoop/sqoop-1.4.5-cdh5.4.0-changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.5+cdh5.4.0+26	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.5-cdh5.4.0.tar.gz	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.5-cdh5.4.0-release.html	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.5-cdh5.4.0-changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.4.0+13	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.4.0.tar.gz	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.4.0-release.html	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.4.0-changes.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.4.0+84	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.4.0.tar.gz	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.4.0-release.html	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.4.0-changes.txt
		Tarball	Release notes	Changes

CDH 5.3.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.3.x release are listed in the following table. For the same type of information for other CDH releases, see [Packaging and Tarballs](#). To view the overall release notes for CDH 5.x.x, see the [Release Notes](#).



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.3.10 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.10+99	https://archive.apache.org/dist/cdh/5/avro-1.7.6+cdh5.3.10.tar.gz	https://archive.apache.org/dist/cdh/5/avro-1.7.6+cdh5.3.10-release.html	https://archive.apache.org/dist/cdh/5/avro-1.7.6+cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.10+48	https://archive.apache.org/dist/cdh/5/crunch-0.11.0+cdh5.3.10.tar.gz	https://archive.apache.org/dist/cdh/5/crunch-0.11.0+cdh5.3.10-release.html	https://archive.apache.org/dist/cdh/5/crunch-0.11.0+cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Datafu	pigudatafu-1.10+cdh5.3.10+24	https://archive.apache.org/dist/cdh/5/datafu-1.10+cdh5.3.10.tar.gz	https://archive.apache.org/dist/cdh/5/datafu-1.10+cdh5.3.10-release.html	https://archive.apache.org/dist/cdh/5/datafu-1.10+cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Flume-ng	flume-ng-1.5.0+cdh5.3.10+102	https://archive.apache.org/dist/cdh/5/flume-ng-1.5.0+cdh5.3.10.tar.gz	https://archive.apache.org/dist/cdh/5/flume-ng-1.5.0+cdh5.3.10-release.html	https://archive.apache.org/dist/cdh/5/flume-ng-1.5.0+cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.10+1001	https://archive.apache.org/dist/cdh/5/hadoop-2.5.0+cdh5.3.10.tar.gz	https://archive.apache.org/dist/cdh/5/hadoop-2.5.0+cdh5.3.10-release.html	https://archive.apache.org/dist/cdh/5/hadoop-2.5.0+cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
HBase	hbase-0.98.6+cdh5.3.10+159	https://archive.apache.org/dist/cdh/5/hbase-0.98.6+cdh5.3.10.tar.gz	https://archive.apache.org/dist/cdh/5/hbase-0.98.6+cdh5.3.10-release.html	https://archive.apache.org/dist/cdh/5/hbase-0.98.6+cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
HBase-solr	hbase-solr-1.5+cdh5.3.10+35	https://archive.apache.org/dist/cdh/5/hbase-solr-1.5+cdh5.3.10.tar.gz	https://archive.apache.org/dist/cdh/5/hbase-solr-1.5+cdh5.3.10-release.html	https://archive.apache.org/dist/cdh/5/hbase-solr-1.5+cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.10+467	https://archive.apache.org/dist/cdh/5/hive-0.13.1+cdh5.3.10.tar.gz	https://archive.apache.org/dist/cdh/5/hive-0.13.1+cdh5.3.10-release.html	https://archive.apache.org/dist/cdh/5/hive-0.13.1+cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://s3.amazonaws.com/cdh5/cdh/5/hue-0.13.1-cdh5.3.10.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-0.13.1-cdh5.3.10-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-0.13.1-cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.10+217	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0-cdh5.3.10.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0-cdh5.3.10-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0-cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Apache Impala	impala-2.1.8+cdh5.3.10+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.1.8-cdh5.3.10-release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.1.8-cdh5.3.10-CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.3.10+208	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0-cdh5.3.10.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0-cdh5.3.10-release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0-cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.3.10+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.3.10.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.3.10-release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.3.10+31	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9-cdh5.3.10.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9-cdh5.3.10-release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9-cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.3.10+355	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0-cdh5.3.10.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0-cdh5.3.10-release.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0-cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes
Apache Parquet	parquet-1.5.0+cdh5.3.10+72	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0-cdh5.3.10.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0-cdh5.3.10-release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0-cdh5.3.10-CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Parquet-format	parquet-format-2.10.0+cdh5.3.10+17	https://us.apache.org/dist/mirrors/apache/parquet-format/2.10.0+cdh5.3.10+17/parquet-format-2.10.0+cdh5.3.10+17.tar.gz	https://us.apache.org/dist/mirrors/apache/parquet-format/2.10.0+cdh5.3.10+17/parquet-format-2.10.0+cdh5.3.10+17.html	https://us.apache.org/dist/mirrors/apache/parquet-format/2.10.0+cdh5.3.10+17/parquet-format-2.10.0+cdh5.3.10+17.html
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.3.10+67	https://us.apache.org/dist/mirrors/apache/pig/0.12.0+cdh5.3.10+67/pig-0.12.0+cdh5.3.10+67.tar.gz	https://us.apache.org/dist/mirrors/apache/pig/0.12.0+cdh5.3.10+67/pig-0.12.0+cdh5.3.10+67.html	https://us.apache.org/dist/mirrors/apache/pig/0.12.0+cdh5.3.10+67/pig-0.12.0+cdh5.3.10+67.html
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.3.10+0	https://us.apache.org/dist/mirrors/apache/search/1.0.0+cdh5.3.10+0/search-1.0.0+cdh5.3.10+0.tar.gz	https://us.apache.org/dist/mirrors/apache/search/1.0.0+cdh5.3.10+0/search-1.0.0+cdh5.3.10+0.html	https://us.apache.org/dist/mirrors/apache/search/1.0.0+cdh5.3.10+0/search-1.0.0+cdh5.3.10+0.html
		Tarball	Release notes	Changes
Apache Sentry (Incubating)	sentry-1.4.0+cdh5.3.10+177	https://us.apache.org/dist/mirrors/apache/sentry/1.4.0+cdh5.3.10+177/sentry-1.4.0+cdh5.3.10+177.tar.gz	https://us.apache.org/dist/mirrors/apache/sentry/1.4.0+cdh5.3.10+177/sentry-1.4.0+cdh5.3.10+177.html	https://us.apache.org/dist/mirrors/apache/sentry/1.4.0+cdh5.3.10+177/sentry-1.4.0+cdh5.3.10+177.html
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.10+361	https://us.apache.org/dist/mirrors/apache/solr/4.4.0+cdh5.3.10+361/solr-4.4.0+cdh5.3.10+361.tar.gz	https://us.apache.org/dist/mirrors/apache/solr/4.4.0+cdh5.3.10+361/solr-4.4.0+cdh5.3.10+361.html	https://us.apache.org/dist/mirrors/apache/solr/4.4.0+cdh5.3.10+361/solr-4.4.0+cdh5.3.10+361.html
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.10+389	https://us.apache.org/dist/mirrors/apache/spark/1.2.0+cdh5.3.10+389/spark-1.2.0+cdh5.3.10+389.tar.gz	https://us.apache.org/dist/mirrors/apache/spark/1.2.0+cdh5.3.10+389/spark-1.2.0+cdh5.3.10+389.html	https://us.apache.org/dist/mirrors/apache/spark/1.2.0+cdh5.3.10+389/spark-1.2.0+cdh5.3.10+389.html
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.10+81	https://us.apache.org/dist/mirrors/apache/sqoop/1.4.5+cdh5.3.10+81/sqoop-1.4.5+cdh5.3.10+81.tar.gz	https://us.apache.org/dist/mirrors/apache/sqoop/1.4.5+cdh5.3.10+81/sqoop-1.4.5+cdh5.3.10+81.html	https://us.apache.org/dist/mirrors/apache/sqoop/1.4.5+cdh5.3.10+81/sqoop-1.4.5+cdh5.3.10+81.html
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.9.4+cdh5.3.10+37	https://us.apache.org/dist/mirrors/apache/sqoop2/1.9.4+cdh5.3.10+37/sqoop2-1.9.4+cdh5.3.10+37.tar.gz	https://us.apache.org/dist/mirrors/apache/sqoop2/1.9.4+cdh5.3.10+37/sqoop2-1.9.4+cdh5.3.10+37.html	https://us.apache.org/dist/mirrors/apache/sqoop2/1.9.4+cdh5.3.10+37/sqoop2-1.9.4+cdh5.3.10+37.html
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.3.10+23	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.3.10.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.3.10.release.html	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.3.10.CHANGES.txt
		Tarball	Release notes	Changes
Zookeeper	zookeeper-3.4.5+cdh5.3.10+94	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.10.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.10.release.html	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.10.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.3.9 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.9+98	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.3.9.release.html	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.3.9.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.9+47	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.3.9.release.html	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.3.9.CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pig-df-datafu-1.10+cdh5.3.9+23	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.3.9.release.html	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.3.9.CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.3.9+100	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.9.release.html	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.9.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.9+942	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.9.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.9.CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-0.98.6+cdh5.3.9+133	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.9.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.9.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.9.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.3.9+33	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.9.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.9.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.9.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.9+446	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.9.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.9.html	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.9.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.9+210	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.9.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.9.html	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.9.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.1.7+cdh5.3.9+0	(none)	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.1.7+cdh5.3.9.html	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.1.7+cdh5.3.9.html#changes
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.3.9+207	https://www.cloudera.com/p/cdh5/cdh/5/kite-0.15.0+cdh5.3.9.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/kite-0.15.0+cdh5.3.9.html	https://www.cloudera.com/p/cdh5/cdh/5/kite-0.15.0+cdh5.3.9.html#changes
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.3.9+0	https://www.cloudera.com/p/cdh5/cdh/5/llama-1.0.0+cdh5.3.9.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/llama-1.0.0+cdh5.3.9.html	https://www.cloudera.com/p/cdh5/cdh/5/llama-1.0.0+cdh5.3.9.html#changes
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.3.9+30	https://www.cloudera.com/p/cdh5/cdh/5/mahout-0.9+cdh5.3.9.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/mahout-0.9+cdh5.3.9.html	https://www.cloudera.com/p/cdh5/cdh/5/mahout-0.9+cdh5.3.9.html#changes
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.3.9+353	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.9.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.9.PMS.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.3.9+71	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.9.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.9.PMS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.3.9+16	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.3.9.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.3.9.PMS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.3.9+66	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.3.9.html	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.3.9.PMS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.3.9+0	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.3.9.html	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.3.9.PMS.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.3.9+160	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.4.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.4.0+cdh5.3.9.html	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.4.0+cdh5.3.9.PMS.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.9+360	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.4.0+cdh5.3.9.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.4.0+cdh5.3.9.html	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.4.0+cdh5.3.9.PMS.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.9+387	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ sqoop-1.4.5+cdh5.3.9.tar.gz	5/ sqoop-1.4.5+cdh5.3.9.release.html	5/ sqoop-1.4.5+cdh5.3.9.changes.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.9+80	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.3.9.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.3.9.release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.3.9.changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.4+cdh5.3.9+34	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop2-1.99.4+cdh5.3.9.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop2-1.99.4+cdh5.3.9.release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop2-1.99.4+cdh5.3.9.changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.3.9+22	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.3.9.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.3.9.release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.3.9.changes.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.3.9+93	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.3.9.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.3.9.release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.3.9.changes.txt
		Tarball	Release notes	Changes

CDH 5.3.8 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.8+95	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.3.8.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.3.8.release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.3.8.changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.8+46	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ crunch-0.11.0+cdh5.3.8.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ crunch-0.11.0+cdh5.3.8.release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ crunch-0.11.0+cdh5.3.8.changes.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.3.8+22	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ pigudatafu-1.10+cdh5.3.8.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ pigudatafu-1.10+cdh5.3.8.release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ pigudatafu-1.10+cdh5.3.8.changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://www.cloudera.com/p/cdh5/cdh/5/cdh5.11.0+cdh5.3.8.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/cdh5.11.0+cdh5.3.8.html	https://www.cloudera.com/p/cdh5/cdh/5/cdh5.11.0+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.3.8+98	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.8.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.8.html	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.8+925	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.8.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.8.html	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.3.8+126	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.8.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.8.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.3.8+32	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.8.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.8.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.8+431	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.8.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.8.html	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.8+207	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.8.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.8.html	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.1.6+cdh5.3.8+0	(none)	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.1.6+cdh5.3.8.html	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.1.6+cdh5.3.8.html#changes
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.3.8+159	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.4.0+cdh5.3.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.4.0+cdh5.3.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.4.0+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.8+357	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.3.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.3.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.8+383	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.2.0+cdh5.3.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.2.0+cdh5.3.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.2.0+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.8+79	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.3.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.3.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.4+cdh5.3.8+33	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.4+cdh5.3.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.4+cdh5.3.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.4+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.3.8+20	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.3.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.3.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.3.8.html#changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.3.8+92	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.8.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.8.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.8.html#changes
		Tarball	Release notes	Changes

CDH 5.3.6 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
-----------	-----------------	---------	---------------	--------------

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.6+93	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/avro1.7.6+cdh5.3.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/avro1.7.6+cdh5.3.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/avro1.7.6+cdh5.3.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.6+31	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/crunch0.11.0+cdh5.3.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/crunch0.11.0+cdh5.3.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/crunch0.11.0+cdh5.3.6.PAGE.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.3.6+21	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/datafu1.10+cdh5.3.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/datafu1.10+cdh5.3.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/datafu1.10+cdh5.3.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.3.6+93	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/flume-ng1.5.0+cdh5.3.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/flume-ng1.5.0+cdh5.3.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/flume-ng1.5.0+cdh5.3.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.6+898	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop2.5.0+cdh5.3.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop2.5.0+cdh5.3.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop2.5.0+cdh5.3.6.PAGE.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.3.6+115	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase0.98.6+cdh5.3.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase0.98.6+cdh5.3.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase0.98.6+cdh5.3.6.PAGE.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.3.6+31	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-solr1.5+cdh5.3.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-solr1.5+cdh5.3.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-solr1.5+cdh5.3.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.6+397	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive0.13.1+cdh5.3.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive0.13.1+cdh5.3.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive0.13.1+cdh5.3.6.PAGE.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.6+203	https://s3.amazonaws.com/cdh5/cdh/5/hue370cdh536.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue370cdh536.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue370cdh536.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.1.5+cdh5.3.6+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/imp215cdh536.release.html	https://s3.amazonaws.com/cdh5/cdh/5/imp215cdh536.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.3.6+201	https://s3.amazonaws.com/cdh5/cdh/5/kite0150cdh536.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite0150cdh536.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite0150cdh536.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.3.6+0	https://s3.amazonaws.com/cdh5/cdh/5/llama100cdh536.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama100cdh536.release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama100cdh536.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.3.6+25	https://s3.amazonaws.com/cdh5/cdh/5/mahout09cdh536.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout09cdh536.release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout09cdh536.CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.3.6+349	https://s3.amazonaws.com/cdh5/cdh/5/oozie400cdh536.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie400cdh536.release.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie400cdh536.CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.3.6+69	https://s3.amazonaws.com/cdh5/cdh/5/parquet150cdh536.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet150cdh536.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet150cdh536.CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.10+cdh5.3.6+14	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ pig-0.12.0+cdh5.3.6.tar.gz	5/ pig-0.12.0+cdh5.3.6.html	5/ pig-0.12.0+cdh5.3.6.html
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.3.6+59	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.3.6.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.3.6.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.3.6.html
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.3.6+0	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.3.6.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.3.6.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.3.6.html
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.3.6+155	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.4.0+cdh5.3.6.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.4.0+cdh5.3.6.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.4.0+cdh5.3.6.html
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.6+352	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.3.6.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.3.6.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.3.6.html
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.6+379	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-1.2.0+cdh5.3.6.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-1.2.0+cdh5.3.6.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-1.2.0+cdh5.3.6.html
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.6+78	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.3.6.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.3.6.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.3.6.html
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.4+cdh5.3.6+32	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.4+cdh5.3.6.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.4+cdh5.3.6.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.4+cdh5.3.6.html
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Whirr	whirr-0.9.0+cdh5.3.6+19	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.3.6.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.3.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.3.6.CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.3.6+91	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.6.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.6.release.html	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.6.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.3.5 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.5+90	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.3.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.3.5.release.html	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.3.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.5+21	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.3.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.3.5.release.html	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.3.5.CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.3.5+18	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.3.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.3.5.release.html	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.3.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.3.5+88	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.5.release.html	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.5+873	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.5.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.5.CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.3.5+92			

Component	Package Version	Tarball	Release Notes	Changes File
		https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hbase-0.9.6-cdh5.3.5.tar.gz	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hbase-0.9.6-cdh5.3.5.html	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hbase-0.9.6-cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.3.5+28	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hbase-solr-1.5-cdh5.3.5.tar.gz	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hbase-solr-1.5-cdh5.3.5.html	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hbase-solr-1.5-cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.5+366	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hive-0.13.1-cdh5.3.5.tar.gz	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hive-0.13.1-cdh5.3.5.html	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hive-0.13.1-cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.5+196	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hue-3.7.0-cdh5.3.5.tar.gz	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hue-3.7.0-cdh5.3.5.html	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ hue-3.7.0-cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.1.4+cdh5.3.5+0	(none)	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ impala-2.1.4-cdh5.3.5.html	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ impala-2.1.4-cdh5.3.5.PM5.txt
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.3.5+198	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ kite-0.15.0-cdh5.3.5.tar.gz	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ kite-0.15.0-cdh5.3.5.html	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ kite-0.15.0-cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.3.5+0	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.3.5.tar.gz	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.3.5.html	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.3.5+22	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ mahout-0.9-cdh5.3.5.tar.gz	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ mahout-0.9-cdh5.3.5.html	https:// us.apache.org/dist/techtopology/ p/cdh5/cdh/ 5/ mahout-0.9-cdh5.3.5.PM5.txt
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.3.5+345	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.3.5.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.3.5.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.3.5+66	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.3.5.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.3.5.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.3.5+11	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.3.5.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.3.5.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.3.5+55	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.3.5.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.3.5.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.3.5+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.3.5.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.3.5.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.3.5+150	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.3.5.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.3.5.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.5+344	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.3.5.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.3.5.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.3.5.PM5.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.5+374	https://archive.apache.org/dist/spark/p/cdh5/cdh/	https://archive.apache.org/dist/spark/p/cdh5/cdh/	https://archive.apache.org/dist/spark/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ sqoop-1.20-cdh5.3.5.tar.gz	5/ sqoop-1.20-cdh5.3.5-release.html	5/ sqoop-1.20-cdh5.3.5-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.5+75	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop-1.45-cdh5.3.5.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop-1.45-cdh5.3.5-release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop-1.45-cdh5.3.5-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.4+cdh5.3.5+29	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop2-1.94-cdh5.3.5.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop2-1.94-cdh5.3.5-release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ sqoop2-1.94-cdh5.3.5-CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.3.5+16	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ whirr-0.90-cdh5.3.5.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ whirr-0.90-cdh5.3.5-release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ whirr-0.90-cdh5.3.5-CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.3.5+88	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ zookeeper-3.45-cdh5.3.5.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ zookeeper-3.45-cdh5.3.5-release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ zookeeper-3.45-cdh5.3.5-CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.3.4 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.4+89	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ avro-1.76-cdh5.3.4.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ avro-1.76-cdh5.3.4-release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ avro-1.76-cdh5.3.4-CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.4+20	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ crunch-0.110-cdh5.3.4.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ crunch-0.110-cdh5.3.4-release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ crunch-0.110-cdh5.3.4-CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.3.4+17	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ pigudatafu-1.10-cdh5.3.4.tar.gz	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ pigudatafu-1.10-cdh5.3.4-release.html	https:// archive.apache.org/dist/parcels/ p/cdh5/cdh/ 5/ pigudatafu-1.10-cdh5.3.4-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/cdh-1.10+cdh5.3.4.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/cdh-1.10+cdh5.3.4.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/cdh-1.10+cdh5.3.4.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.3.4+87	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.4.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.4.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.4.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.4+871	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.4.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.4.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.4.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.3.4+91	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.4.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.4.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.4.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.3.4+27	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.4.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.4.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.4.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.4+365	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.4.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.4.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.4.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.4+195	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.4.tar.gz	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.4.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.4.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.1.4+cdh5.3.4+0	(none)	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/impala-2.1.4+cdh5.3.4.html	https://s3.amazonaws.com/cloudera-releases/p/cdh5/cdh/5/impala-2.1.4+cdh5.3.4.html#changes
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.3.4+149	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.4.0+cdh5.3.4.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.4.0+cdh5.3.4.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.4.0+cdh5.3.4.html#st=changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.4+343	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.3.4.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.3.4.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.3.4.html#st=changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.4+373	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.2.0+cdh5.3.4.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.2.0+cdh5.3.4.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.2.0+cdh5.3.4.html#st=changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.4+74	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.3.4.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.3.4.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.5+cdh5.3.4.html#st=changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.4+cdh5.3.4+28	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.4+cdh5.3.4.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.4+cdh5.3.4.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.4+cdh5.3.4.html#st=changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.3.4+15	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.3.4.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.3.4.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.3.4.html#st=changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.3.4+87	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.4.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.4.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.4.html#st=changes
		Tarball	Release notes	Changes

CDH 5.3.3 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
-----------	-----------------	---------	---------------	--------------

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.3+88	https://archive.apache.org/dist/avro/1.7.6+cdh5.3.3+88.tar.gz	https://archive.apache.org/dist/avro/1.7.6+cdh5.3.3+88/RELEASENOTES.html	https://archive.apache.org/dist/avro/1.7.6+cdh5.3.3+88/CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.3+19	https://archive.apache.org/dist/crunch/0.11.0+cdh5.3.3+19.tar.gz	https://archive.apache.org/dist/crunch/0.11.0+cdh5.3.3+19/RELEASENOTES.html	https://archive.apache.org/dist/crunch/0.11.0+cdh5.3.3+19/CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.3.3+16	https://archive.apache.org/dist/datafu/1.10+cdh5.3.3+16.tar.gz	https://archive.apache.org/dist/datafu/1.10+cdh5.3.3+16/RELEASENOTES.html	https://archive.apache.org/dist/datafu/1.10+cdh5.3.3+16/CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.3.3+86	https://archive.apache.org/dist/flume-ng/1.5.0+cdh5.3.3+86.tar.gz	https://archive.apache.org/dist/flume-ng/1.5.0+cdh5.3.3+86/RELEASENOTES.html	https://archive.apache.org/dist/flume-ng/1.5.0+cdh5.3.3+86/CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.3+844	https://archive.apache.org/dist/hadoop-core/2.5.0+cdh5.3.3+844.tar.gz	https://archive.apache.org/dist/hadoop-core/2.5.0+cdh5.3.3+844/RELEASENOTES.html	https://archive.apache.org/dist/hadoop-core/2.5.0+cdh5.3.3+844/CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.3.3+85	https://archive.apache.org/dist/hbase/0.98.6+cdh5.3.3+85.tar.gz	https://archive.apache.org/dist/hbase/0.98.6+cdh5.3.3+85/RELEASENOTES.html	https://archive.apache.org/dist/hbase/0.98.6+cdh5.3.3+85/CHANGES.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.3.3+26	https://archive.apache.org/dist/hbase-solr/1.5+cdh5.3.3+26.tar.gz	https://archive.apache.org/dist/hbase-solr/1.5+cdh5.3.3+26/RELEASENOTES.html	https://archive.apache.org/dist/hbase-solr/1.5+cdh5.3.3+26/CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.3+350	https://archive.apache.org/dist/hive/0.13.1+cdh5.3.3+350.tar.gz	https://archive.apache.org/dist/hive/0.13.1+cdh5.3.3+350/RELEASENOTES.html	https://archive.apache.org/dist/hive/0.13.1+cdh5.3.3+350/CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.3+180	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0+cdh5.3.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0+cdh5.3.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0+cdh5.3.3.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.1.3+cdh5.3.3+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.1.3+cdh5.3.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.1.3+cdh5.3.3.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.3.3+194	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.3.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.3.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.3.3.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.3.3+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.3.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.3.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.3.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.3.3+20	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.3.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.3.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.3.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.3.3+340	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.3.CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.3.3+64	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.3.CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.3.3+9	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ pig-0.12.0+cdh5.3.3+53.tar.gz	5/ pig-0.12.0+cdh5.3.3+53.html	5/ pig-0.12.0+cdh5.3.3+53.html
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.3.3+53	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.3.3+53.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.3.3+53.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.3.3+53.html
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.3.3+0	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.3.3+0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.3.3+0.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.3.3+0.html
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.3.3+137	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.4.0+cdh5.3.3+137.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.4.0+cdh5.3.3+137.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.4.0+cdh5.3.3+137.html
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.3+337	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.3.3+337.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.3.3+337.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.3.3+337.html
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.3+371	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-1.2.0+cdh5.3.3+371.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-1.2.0+cdh5.3.3+371.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-1.2.0+cdh5.3.3+371.html
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.3+67	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.3.3+67.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.3.3+67.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.3.3+67.html
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.4+cdh5.3.3+23	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.4+cdh5.3.3+23.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.4+cdh5.3.3+23.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.4+cdh5.3.3+23.html
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Whirr	whirr-0.9.0+cdh5.3.3+14	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.3.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.3.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.3.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.3.3+84	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.3.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.3.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.2+87	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.6+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.2+18	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.11.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.3.2+14	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.3.2+84	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.2+813	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.3.2+83			

Component	Package Version	Tarball	Release Notes	Changes File
		https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.96-cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.96-cdh5.3.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.96-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.3.2+25	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.5-cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.5-cdh5.3.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.5-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.2+330	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.13.1-cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.13.1-cdh5.3.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.13.1-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.2+163	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0-cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0-cdh5.3.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.1.2+cdh5.3.2+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.1.2-cdh5.3.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.1.2-cdh5.3.2-CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.3.2+193	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0-cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0-cdh5.3.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.3.2+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.3.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.3.2+19	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9-cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9-cdh5.3.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9-cdh5.3.2-CHANGES.txt
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.3.2+339	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.3.2+62	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.3.2+8	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-2.1.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.3.2+51	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.3.2+0	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.3.2+128	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.4.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.4.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.4.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.2+326	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.4.0+cdh5.3.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.4.0+cdh5.3.2.release.html	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.4.0+cdh5.3.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.2+369	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ parquet-1.2.0-cdh5.3.2.tar.gz	5/ parquet-1.2.0-cdh5.3.2-release.html	5/ parquet-1.2.0-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.2+64	https://archive.apache.org/dist/sqoop/sqoop-1.4.5-cdh5.3.2.tar.gz	https://archive.apache.org/dist/sqoop/sqoop-1.4.5-cdh5.3.2-release.html	https://archive.apache.org/dist/sqoop/sqoop-1.4.5-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.4+cdh5.3.2+21	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.4-cdh5.3.2.tar.gz	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.4-cdh5.3.2-release.html	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.4-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.3.2+13	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.3.2.tar.gz	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.3.2-release.html	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.3.2+83	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.3.2.tar.gz	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.3.2-release.html	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.3.2-CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.3.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.1+74	https://archive.apache.org/dist/avro/avro-1.7.6-cdh5.3.1.tar.gz	https://archive.apache.org/dist/avro/avro-1.7.6-cdh5.3.1-release.html	https://archive.apache.org/dist/avro/avro-1.7.6-cdh5.3.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.1+17	https://archive.apache.org/dist/crunch/crunch-0.11.0-cdh5.3.1.tar.gz	https://archive.apache.org/dist/crunch/crunch-0.11.0-cdh5.3.1-release.html	https://archive.apache.org/dist/crunch/crunch-0.11.0-cdh5.3.1-CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.3.1+13	https://archive.apache.org/dist/pigudatafu/pigudatafu-1.10-cdh5.3.1.tar.gz	https://archive.apache.org/dist/pigudatafu/pigudatafu-1.10-cdh5.3.1-release.html	https://archive.apache.org/dist/pigudatafu/pigudatafu-1.10-cdh5.3.1-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://www.cloudera.com/p/cdh5/cdh/5/cdh5.1.0+cdh5.3.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/cdh5.1.0+cdh5.3.1.html	https://www.cloudera.com/p/cdh5/cdh/5/cdh5.1.0+cdh5.3.1.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.3.1+80	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.1.html	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.3.1.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.1+791	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.3.1.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.3.1+74	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.3.1.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.3.1+24	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.3.1.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.1+308	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.3.1.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.1+135	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.7.0+cdh5.3.1.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.1.1+cdh5.3.1+0	(none)	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.1.1+cdh5.3.1.html	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.1.1+cdh5.3.1.html#changes
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.3.1+127	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.4.0+cdh5.3.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.4.0+cdh5.3.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.4.0+cdh5.3.1Changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.1+315	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.4.0+cdh5.3.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.4.0+cdh5.3.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.4.0+cdh5.3.1Changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.1+365	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.2.0+cdh5.3.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.2.0+cdh5.3.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.2.0+cdh5.3.1Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.1+61	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.3.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.3.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.3.1Changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.4+cdh5.3.1+20	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.4+cdh5.3.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.4+cdh5.3.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.4+cdh5.3.1Changes.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.3.1+12	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.3.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.3.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.3.1Changes.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.3.1+82	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.1.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.1.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.3.1Changes.txt
		Tarball	Release notes	Changes

CDH 5.3.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
-----------	-----------------	---------	---------------	--------------

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.3.0+73	https://archive.apache.org/dist/avro/1.7.6+cdh5.3.0+73/apache-avro-1.7.6+cdh5.3.0+73.tar.gz	https://archive.apache.org/dist/avro/1.7.6+cdh5.3.0+73/avro-1.7.6+cdh5.3.0+73-releasenotes.html	https://archive.apache.org/dist/avro/1.7.6+cdh5.3.0+73/avro-1.7.6+cdh5.3.0+73-changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.3.0+16	https://archive.apache.org/dist/crunch/0.11.0+cdh5.3.0+16/apache-crunch-0.11.0+cdh5.3.0+16.tar.gz	https://archive.apache.org/dist/crunch/0.11.0+cdh5.3.0+16/crunch-0.11.0+cdh5.3.0+16-releasenotes.html	https://archive.apache.org/dist/crunch/0.11.0+cdh5.3.0+16/crunch-0.11.0+cdh5.3.0+16-changes.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.3.0+12	https://archive.apache.org/dist/datafu/1.10+cdh5.3.0+12/datafu-1.10+cdh5.3.0+12.tar.gz	https://archive.apache.org/dist/datafu/1.10+cdh5.3.0+12/datafu-1.10+cdh5.3.0+12-releasenotes.html	https://archive.apache.org/dist/datafu/1.10+cdh5.3.0+12/datafu-1.10+cdh5.3.0+12-changes.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.3.0+79	https://archive.apache.org/dist/flume-ng/1.5.0+cdh5.3.0+79/flume-ng-1.5.0+cdh5.3.0+79.tar.gz	https://archive.apache.org/dist/flume-ng/1.5.0+cdh5.3.0+79/flume-ng-1.5.0+cdh5.3.0+79-releasenotes.html	https://archive.apache.org/dist/flume-ng/1.5.0+cdh5.3.0+79/flume-ng-1.5.0+cdh5.3.0+79-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.3.0+781	https://archive.apache.org/dist/hadoop-core/2.5.0+cdh5.3.0+781/hadoop-2.5.0+cdh5.3.0+781.tar.gz	https://archive.apache.org/dist/hadoop-core/2.5.0+cdh5.3.0+781/hadoop-2.5.0+cdh5.3.0+781-releasenotes.html	https://archive.apache.org/dist/hadoop-core/2.5.0+cdh5.3.0+781/hadoop-2.5.0+cdh5.3.0+781-changes.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.3.0+73	https://archive.apache.org/dist/hbase/0.98.6+cdh5.3.0+73/hbase-0.98.6+cdh5.3.0+73.tar.gz	https://archive.apache.org/dist/hbase/0.98.6+cdh5.3.0+73/hbase-0.98.6+cdh5.3.0+73-releasenotes.html	https://archive.apache.org/dist/hbase/0.98.6+cdh5.3.0+73/hbase-0.98.6+cdh5.3.0+73-changes.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.3.0+23	https://archive.apache.org/dist/hbase-solr/1.5+cdh5.3.0+23/hbase-solr-1.5+cdh5.3.0+23.tar.gz	https://archive.apache.org/dist/hbase-solr/1.5+cdh5.3.0+23/hbase-solr-1.5+cdh5.3.0+23-releasenotes.html	https://archive.apache.org/dist/hbase-solr/1.5+cdh5.3.0+23/hbase-solr-1.5+cdh5.3.0+23-changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.3.0+306	https://archive.apache.org/dist/hive/0.13.1+cdh5.3.0+306/hive-0.13.1+cdh5.3.0+306.tar.gz	https://archive.apache.org/dist/hive/0.13.1+cdh5.3.0+306/hive-0.13.1+cdh5.3.0+306-releasenotes.html	https://archive.apache.org/dist/hive/0.13.1+cdh5.3.0+306/hive-0.13.1+cdh5.3.0+306-changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-3.7.0+cdh5.3.0+134	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0+cdh5.3.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0+cdh5.3.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.7.0+cdh5.3.0.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.1+cdh5.3.0+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.1.0+cdh5.3.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.1.0+cdh5.3.0.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.3.0+190	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.3.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.3.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.3.0.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.3.0+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.3.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.3.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.3.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.3.0+17	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.3.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.3.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.3.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.3.0+334	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.3.0.CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.3.0+52	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.0.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.3.0.CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.3.0+6	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ parquet-1.10.0-cdh5.3.0.tar.gz	5/ parquet-1.10.0-cdh5.3.0.html	5/ parquet-1.10.0-cdh5.3.0.html#changes
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.3.0+46	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0-cdh5.3.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0-cdh5.3.0.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0-cdh5.3.0.html#changes
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.3.0+0	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0-cdh5.3.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0-cdh5.3.0.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0-cdh5.3.0.html#changes
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.3.0+126	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.4.0-cdh5.3.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.4.0-cdh5.3.0.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.4.0-cdh5.3.0.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.3.0+314	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0-cdh5.3.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0-cdh5.3.0.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0-cdh5.3.0.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.2.0+cdh5.3.0+364	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-1.2.0-cdh5.3.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-1.2.0-cdh5.3.0.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-1.2.0-cdh5.3.0.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.3.0+60	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.5-cdh5.3.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.5-cdh5.3.0.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.5-cdh5.3.0.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.4+cdh5.3.0+19	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.4-cdh5.3.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.4-cdh5.3.0.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.4-cdh5.3.0.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Whirr	whirr-0.9.0+cdh5.3.0+11	https://archive.apache.org/dist/whirr/whirr-0.9.0+cdh5.3.0.tar.gz	https://archive.apache.org/dist/whirr/whirr-0.9.0+cdh5.3.0.html	https://archive.apache.org/dist/whirr/whirr-0.9.0+cdh5.3.0/changes.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.3.0+81	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5+cdh5.3.0.tar.gz	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5+cdh5.3.0.html	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5+cdh5.3.0/changes.txt
		Tarball	Release notes	Changes

CDH 5.2.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.2.x release are listed in the following table. For the same type of information for other CDH releases, see [Packaging and Tarballs](#). To view the overall release notes for CDH 5.x.x, see the [Release Notes](#).



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.2.6 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.2.6+81	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/avro-1.7.6+cdh5.2.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/avro-1.7.6+cdh5.2.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/avro-1.7.6+cdh5.2.6.PMS.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.2.6+31	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/crunch-0.11.0+cdh5.2.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/crunch-0.11.0+cdh5.2.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/crunch-0.11.0+cdh5.2.6.PMS.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.2.6+22	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/datafu-1.10+cdh5.2.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/datafu-1.10+cdh5.2.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/datafu-1.10+cdh5.2.6.PMS.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.2.6+67	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.2.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.2.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.2.6.PMS.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.2.6+72	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.2.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.2.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.2.6.PMS.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.2.6+99	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-0.98.6+cdh5.2.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-0.98.6+cdh5.2.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-0.98.6+cdh5.2.6.PMS.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.2.6+27	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.2.6.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.2.6.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.2.6.PMS.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.2.6+281	https://us.apache.org/dist/hadoop/p/cdh5/cdh/	https://us.apache.org/dist/hadoop/p/cdh5/cdh/	https://us.apache.org/dist/hadoop/p/cdh5/cdh/

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		5/ hive-0.13.1-cdh5.2.6.tar.gz	5/ hive-0.13.1-cdh5.2.6-releasenotes.html	5/ hive-0.13.1-cdh5.2.6-changes.txt
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.2.6+558	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ hue-3.6.0-cdh5.2.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ hue-3.6.0-cdh5.2.6-releasenotes.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ hue-3.6.0-cdh5.2.6-changes.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.0.5+cdh5.2.6+0	(none)	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ impala-2.0.5-cdh5.2.6-releasenotes.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ impala-2.0.5-cdh5.2.6-changes.txt
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.2.6+161	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ kite-0.15.0-cdh5.2.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ kite-0.15.0-cdh5.2.6-releasenotes.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ kite-0.15.0-cdh5.2.6-changes.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.2.6+0	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.2.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.2.6-releasenotes.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.2.6-changes.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.2.6+25	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ mahout-0.9-cdh5.2.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ mahout-0.9-cdh5.2.6-releasenotes.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ mahout-0.9-cdh5.2.6-changes.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.2.6+335	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ oozie-4.0.0-cdh5.2.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ oozie-4.0.0-cdh5.2.6-releasenotes.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ oozie-4.0.0-cdh5.2.6-changes.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.2.6+46	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ parquet-1.5.0-cdh5.2.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ parquet-1.5.0-cdh5.2.6-releasenotes.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ parquet-1.5.0-cdh5.2.6-changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Parquet-format	parquet-format-2.10+cdh5.2.6+12	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ parquet-format-2.10+cdh5.2.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ parquet-format-2.10+cdh5.2.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ parquet-format-2.10+cdh5.2.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.2.6+53	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.2.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.2.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.2.6.PAGE.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.2.6+0	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.2.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.2.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.2.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.2.6+97	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sentry-1.4.0+cdh5.2.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sentry-1.4.0+cdh5.2.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sentry-1.4.0+cdh5.2.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.2.6+315	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.2.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.2.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.2.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.1.0+cdh5.2.6+71	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-1.1.0+cdh5.2.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-1.1.0+cdh5.2.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-1.1.0+cdh5.2.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.2.6+57	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.2.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.2.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.2.6.PAGE.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.2.6+38	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.2.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.2.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.2.6.PAGE.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.2.6+18	https://archive.apache.org/dist/whirr/whirr-0.9.0+cdh5.2.6.tar.gz	https://archive.apache.org/dist/whirr/whirr-0.9.0+cdh5.2.6.html	https://archive.apache.org/dist/whirr/whirr-0.9.0+cdh5.2.6.html#changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.2.6+92	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5+cdh5.2.6.tar.gz	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5+cdh5.2.6.html	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5+cdh5.2.6.html#changes
		Tarball	Release notes	Changes

CDH 5.2.5 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.2.5+80	https://archive.apache.org/dist/avro/avro-1.7.6+cdh5.2.5.tar.gz	https://archive.apache.org/dist/avro/avro-1.7.6+cdh5.2.5.html	https://archive.apache.org/dist/avro/avro-1.7.6+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.2.5+24	https://archive.apache.org/dist/crunch/crunch-0.11.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/crunch/crunch-0.11.0+cdh5.2.5.html	https://archive.apache.org/dist/crunch/crunch-0.11.0+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
DataFu	pig-df-datafu-1.10+cdh5.2.5+21	https://archive.apache.org/dist/datafu/pig-df-datafu-1.10+cdh5.2.5.tar.gz	https://archive.apache.org/dist/datafu/pig-df-datafu-1.10+cdh5.2.5.html	https://archive.apache.org/dist/datafu/pig-df-datafu-1.10+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.2.5+65	https://archive.apache.org/dist/flume/flume-ng-1.5.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/flume/flume-ng-1.5.0+cdh5.2.5.html	https://archive.apache.org/dist/flume/flume-ng-1.5.0+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.2.5+644	https://archive.apache.org/dist/hadoop/core/hadoop-2.5.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.5.0+cdh5.2.5.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.5.0+cdh5.2.5.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-0.98.6+cdh5.2.5+78	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.2.5.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.2.5.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.2.5+26	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.2.5.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.2.5.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.2.5+276	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.2.5.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.2.5.html	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.2.5+556	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.6.0+cdh5.2.5.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.6.0+cdh5.2.5.html	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.6.0+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.0.4+cdh5.2.5+0	(none)	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.0.4+cdh5.2.5.html	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.0.4+cdh5.2.5.html#changes
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.2.5+160	https://www.cloudera.com/p/cdh5/cdh/5/kite-0.15.0+cdh5.2.5.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/kite-0.15.0+cdh5.2.5.html	https://www.cloudera.com/p/cdh5/cdh/5/kite-0.15.0+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.2.5+0	https://www.cloudera.com/p/cdh5/cdh/5/llama-1.0.0+cdh5.2.5.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/llama-1.0.0+cdh5.2.5.html	https://www.cloudera.com/p/cdh5/cdh/5/llama-1.0.0+cdh5.2.5.html#changes
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.2.5+24	https://www.cloudera.com/p/cdh5/cdh/5/mahout-0.9+cdh5.2.5.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/mahout-0.9+cdh5.2.5.html	https://www.cloudera.com/p/cdh5/cdh/5/mahout-0.9+cdh5.2.5.html#changes
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.2.5+332	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.2.5.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.2.5.PAGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.2.5+45	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.2.5.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.5.0+cdh5.2.5.PAGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.2.5+11	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.2.5.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-2.1.0+cdh5.2.5.PAGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.2.5+52	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.2.5.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.2.5.PAGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.2.5+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.2.5.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.2.5.PAGES.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.2.5+94	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.2.5.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.4.0+cdh5.2.5.PAGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.2.5+310	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.2.5.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.2.5.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.2.5.PAGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.1.0+cdh5.2.5+70	https://archive.apache.org/dist/spark/p/cdh5/cdh/	https://archive.apache.org/dist/spark/p/cdh5/cdh/	https://archive.apache.org/dist/spark/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ sqoop-1.10+cdh5.2.5.tar.gz	5/ sqoop-1.10+cdh5.2.5.html	5/ sqoop-1.10+cdh5.2.5.html.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.2.5+56	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.2.5.tar.gz	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.2.5.html	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ sqoop-1.4.5+cdh5.2.5.html.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.2.5+37	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.2.5.tar.gz	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.2.5.html	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.2.5.html.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.2.5+17	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.2.5.tar.gz	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.2.5.html	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.2.5.html.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.2.5+89	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.2.5.tar.gz	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.2.5.html	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.2.5.html.txt
		Tarball	Release notes	Changes

CDH 5.2.4 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.2.4+79	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.2.4.tar.gz	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.2.4.html	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ avro-1.7.6+cdh5.2.4.html.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.2.4+22	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ crunch-0.11.0+cdh5.2.4.tar.gz	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ crunch-0.11.0+cdh5.2.4.html	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ crunch-0.11.0+cdh5.2.4.html.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.2.4+20	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ pigudatafu-1.10+cdh5.2.4.tar.gz	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ pigudatafu-1.10+cdh5.2.4.html	https:// archive.apache.org/dist/ parcels/p/cdh5/cdh/ 5/ pigudatafu-1.10+cdh5.2.4.html.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://www.cloudera.com/p/cdh5/cdh/5/cdh110-cdh524.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/cdh110-cdh524-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/cdh110-cdh524-changes.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.2.4+63	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0-cdh524.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0-cdh524-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0-cdh524-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.2.4+618	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0-cdh524.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0-cdh524-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0-cdh524-changes.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.2.4+77	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6-cdh524.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6-cdh524-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6-cdh524-changes.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.2.4+25	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5-cdh524.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5-cdh524-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5-cdh524-changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.2.4+260	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1-cdh524.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1-cdh524-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1-cdh524-changes.txt
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.2.4+550	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.6.0-cdh524.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.6.0-cdh524-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.6.0-cdh524-changes.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.0.3+cdh5.2.4+0	(none)	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.0.3-cdh524-releasenotes.html	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.0.3-cdh524-changes.txt
			Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Kite SDK	kite-0.15.0+cdh5.2.4+159	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.2.4.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.2.4.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.2.4.changes.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.2.4+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.2.4.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.2.4.release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.2.4.changes.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.2.4+23	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.2.4.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.2.4.release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.2.4.changes.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.2.4+331	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.2.4.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.2.4.release.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.2.4.changes.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.2.4+44	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.2.4.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.2.4.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.2.4.changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.2.10+cdh5.2.4+10	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.2.10+cdh5.2.4.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.2.10+cdh5.2.4.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.2.10+cdh5.2.4.changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.2.4+51	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.2.4.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.2.4.release.html	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.2.4.changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.2.4+0	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.2.4.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.2.4.release.html	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.2.4.changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.2.4+92	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.4.0+cdh5.2.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.4.0+cdh5.2.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.4.0+cdh5.2.4.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.2.4+299	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.4.0+cdh5.2.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.4.0+cdh5.2.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.4.0+cdh5.2.4.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.1.0+cdh5.2.4+69	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.1.0+cdh5.2.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.1.0+cdh5.2.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.1.0+cdh5.2.4.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.2.4+53	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.2.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.2.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.2.4.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.2.4+36	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.3+cdh5.2.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.3+cdh5.2.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.3+cdh5.2.4.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.2.4+16	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.2.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.2.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.2.4.html#changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.2.4+88	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.2.4.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.2.4.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.2.4.html#changes
		Tarball	Release notes	Changes

CDH 5.2.3 Packaging and Tarballs



Note: There is no CDH 5.0.7, 5.1.1, 5.1.6, 5.2.2, 5.2.7, 5.3.7, or 5.4.6 release.

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.2.3+78	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/avro-1.7.6+cdh5.2.3.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/avro-1.7.6+cdh5.2.3.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/avro-1.7.6+cdh5.2.3.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.2.3+21	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/crunch-0.11.0+cdh5.2.3.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/crunch-0.11.0+cdh5.2.3.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/crunch-0.11.0+cdh5.2.3.html#changes
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.2.3+18	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/datafu-1.10+cdh5.2.3.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/datafu-1.10+cdh5.2.3.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/datafu-1.10+cdh5.2.3.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.2.3+61	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.2.3.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.2.3.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.2.3.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.2.3+615	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.2.3.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.2.3.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.2.3.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.2.3+76	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-0.98.6+cdh5.2.3.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-0.98.6+cdh5.2.3.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-0.98.6+cdh5.2.3.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.2.3+24	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.2.3.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.2.3.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.2.3.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.2.3+250	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-0.13.1+cdh5.2.3.tar.gz	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-0.13.1+cdh5.2.3.html	https://us.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-0.13.1+cdh5.2.3.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.2.3+548	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.6.0+cdh5.2.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.6.0+cdh5.2.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.6.0+cdh5.2.3.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.0.2+cdh5.2.3+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.0.2+cdh5.2.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.0.2+cdh5.2.3.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.2.3+158	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.2.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.2.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0+cdh5.2.3.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.2.3+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.2.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.2.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.2.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.2.3+22	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.2.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.2.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9+cdh5.2.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.2.3+330	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.2.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.2.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.2.3.CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.5.0+cdh5.2.3+43	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.2.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.2.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.5.0+cdh5.2.3.CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-2.1.0+cdh5.2.3+9	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		5/ parcels-210-cdh5.2.3.tar.gz	5/ parcels-210-cdh5.2.3.html	5/ parcels-210-cdh5.2.3.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.2.3+51	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ pig-0.12.0-cdh5.2.3.tar.gz	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ pig-0.12.0-cdh5.2.3.html	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ pig-0.12.0-cdh5.2.3.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.2.3+0	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ search-1.0.0-cdh5.2.3.tar.gz	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ search-1.0.0-cdh5.2.3.html	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ search-1.0.0-cdh5.2.3.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.2.3+91	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ sentry-1.4.0-cdh5.2.3.tar.gz	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ sentry-1.4.0-cdh5.2.3.html	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ sentry-1.4.0-cdh5.2.3.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.2.3+290	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ solr-4.4.0-cdh5.2.3.tar.gz	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ solr-4.4.0-cdh5.2.3.html	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ solr-4.4.0-cdh5.2.3.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.1.0+cdh5.2.3+67	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ spark-1.1.0-cdh5.2.3.tar.gz	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ spark-1.1.0-cdh5.2.3.html	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ spark-1.1.0-cdh5.2.3.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.2.3+52	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ sqoop-1.4.5-cdh5.2.3.tar.gz	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ sqoop-1.4.5-cdh5.2.3.html	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ sqoop-1.4.5-cdh5.2.3.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.2.3+35	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ sqoop2-1.99.3-cdh5.2.3.tar.gz	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ sqoop2-1.99.3-cdh5.2.3.html	https:// archive.apache.org/dist/ parcels/210/cdh5/cdh/ 5/ sqoop2-1.99.3-cdh5.2.3.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Whirr	whirr-0.9.0+cdh5.2.3+15	https://archive.apache.org/dist/whirr/p/cdh5/cdh/5/whirr-0.9.0+cdh5.2.3.tar.gz	https://archive.apache.org/dist/whirr/p/cdh5/cdh/5/whirr-0.9.0+cdh5.2.3.html	https://archive.apache.org/dist/whirr/p/cdh5/cdh/5/whirr-0.9.0+cdh5.2.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.2.3+87	https://archive.apache.org/dist/zookeeper/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.2.3.tar.gz	https://archive.apache.org/dist/zookeeper/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.2.3.html	https://archive.apache.org/dist/zookeeper/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.2.3.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.2.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.2.1+69	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.6+cdh5.2.1.tar.gz	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.6+cdh5.2.1.html	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.6+cdh5.2.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.2.1+11	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.11.0+cdh5.2.1.tar.gz	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.11.0+cdh5.2.1.html	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.11.0+cdh5.2.1.CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.2.1+15	https://archive.apache.org/dist/datafu/p/cdh5/cdh/5/datafu-1.10+cdh5.2.1.tar.gz	https://archive.apache.org/dist/datafu/p/cdh5/cdh/5/datafu-1.10+cdh5.2.1.html	https://archive.apache.org/dist/datafu/p/cdh5/cdh/5/datafu-1.10+cdh5.2.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.2.1+54	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.2.1.tar.gz	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.2.1.html	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.2.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.2.1+578	https://archive.apache.org/dist/hadoop-core/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.2.1.tar.gz	https://archive.apache.org/dist/hadoop-core/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.2.1.html	https://archive.apache.org/dist/hadoop-core/p/cdh5/cdh/5/hadoop-2.5.0+cdh5.2.1.CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.2.1+64			

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.96-cdh5.21.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.96-cdh5.21-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.96-cdh5.21-CHANGES.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.2.1+21	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.5-cdh5.21.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.5-cdh5.21-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.5-cdh5.21-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.2.1+235	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.13.1-cdh5.21.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.13.1-cdh5.21-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.13.1-cdh5.21-CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.2.1+512	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.6.0-cdh5.21.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.6.0-cdh5.21-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.6.0-cdh5.21-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.0.1+cdh5.2.1+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.0.1-cdh5.21-release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-2.0.1-cdh5.21-CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.15.0+cdh5.2.1+154	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0-cdh5.21.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0-cdh5.21-release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.15.0-cdh5.21-CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.2.1+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.21.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.21-release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.21-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.2.1+19	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9-cdh5.21.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9-cdh5.21-release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.9-cdh5.21-CHANGES.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		5/ sqoop-1.4.5+cdh5.2.1.tar.gz	5/ sqoop-1.4.5+cdh5.2.1-release.html	5/ sqoop-1.4.5+cdh5.2.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.2.1+49	https://archive.apache.org/dist/parcels/parcels/5/sqoop-1.4.5+cdh5.2.1.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/sqoop-1.4.5+cdh5.2.1-release.html	https://archive.apache.org/dist/parcels/parcels/5/sqoop-1.4.5+cdh5.2.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.2.1+32	https://archive.apache.org/dist/parcels/parcels/5/sqoop2-1.99.3+cdh5.2.1.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/sqoop2-1.99.3+cdh5.2.1-release.html	https://archive.apache.org/dist/parcels/parcels/5/sqoop2-1.99.3+cdh5.2.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.2.1+12	https://archive.apache.org/dist/parcels/parcels/5/whirr-0.9.0+cdh5.2.1.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/whirr-0.9.0+cdh5.2.1-release.html	https://archive.apache.org/dist/parcels/parcels/5/whirr-0.9.0+cdh5.2.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.2.1+84	https://archive.apache.org/dist/parcels/parcels/5/zookeeper-3.4.5+cdh5.2.1.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/zookeeper-3.4.5+cdh5.2.1-release.html	https://archive.apache.org/dist/parcels/parcels/5/zookeeper-3.4.5+cdh5.2.1-CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.2.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.6+cdh5.2.0+67	https://archive.apache.org/dist/parcels/parcels/5/avro-1.7.6+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/avro-1.7.6+cdh5.2.0-release.html	https://archive.apache.org/dist/parcels/parcels/5/avro-1.7.6+cdh5.2.0-CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.11.0+cdh5.2.0+9	https://archive.apache.org/dist/parcels/parcels/5/crunch-0.11.0+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/crunch-0.11.0+cdh5.2.0-release.html	https://archive.apache.org/dist/parcels/parcels/5/crunch-0.11.0+cdh5.2.0-CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.2.0+13	https://archive.apache.org/dist/parcels/parcels/5/pigudatafu-1.10+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/pigudatafu-1.10+cdh5.2.0-release.html	https://archive.apache.org/dist/parcels/parcels/5/pigudatafu-1.10+cdh5.2.0-CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		https://www.cloudera.com/p/cdh5/cdh/5/cdh5.11.0-cdh5.2.0.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/cdh5.11.0-cdh5.2.0.html	https://www.cloudera.com/p/cdh5/cdh/5/cdh5.11.0-cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.2.0+49	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0-cdh5.2.0.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0-cdh5.2.0.html	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.5.0-cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.5.0+cdh5.2.0+551	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0-cdh5.2.0.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0-cdh5.2.0.html	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.5.0-cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.6+cdh5.2.0+55	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6-cdh5.2.0.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6-cdh5.2.0.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.98.6-cdh5.2.0.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.2.0+19	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.2.0.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.2.0.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.13.1+cdh5.2.0+221	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1-cdh5.2.0.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1-cdh5.2.0.html	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.13.1-cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.2.0+509	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.6.0-cdh5.2.0.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.6.0-cdh5.2.0.html	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.6.0-cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-2.0+cdh5.2.0+0	(none)	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.0-cdh5.2.0.html	https://www.cloudera.com/p/cdh5/cdh/5/impala-2.0-cdh5.2.0.html#changes
			Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.4.0+cdh5.2.0+84	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.4.0+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.4.0+cdh5.2.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sentry-1.4.0+cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.2.0+282	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.4.0+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.4.0+cdh5.2.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/solr-4.4.0+cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-1.1.0+cdh5.2.0+56	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.1.0+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.1.0+cdh5.2.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/spark-1.1.0+cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.5+cdh5.2.0+47	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.2.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop-1.4.5+cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.2.0+30	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.3+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.3+cdh5.2.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/sqoop2-1.99.3+cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.2.0+10	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.2.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/whirr-0.9.0+cdh5.2.0.html#changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.2.0+82	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.2.0.tar.gz	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.2.0.html	https://archive.apache.org/dist/parcels/cdh5/cdh/5/zookeeper-3.4.5+cdh5.2.0.html#changes
		Tarball	Release notes	Changes

CDH 5.1.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.1.x release are listed in the following table. For

Version and Download Information

the same type of information for other CDH releases, see [Packaging and Tarballs](#). To view the overall release notes for CDH 5.x.x, see the [Release Notes](#).



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.1.5 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.1.5+36	https://archive.apache.org/dist/avro/1.7.5/avro-1.7.5-cdh5.1.5.tar.gz	https://archive.apache.org/dist/avro/1.7.5/avro-1.7.5-cdh5.1.5.html	https://archive.apache.org/dist/avro/1.7.5/avro-1.7.5-cdh5.1.5Changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.10.0+cdh5.1.5+20	https://archive.apache.org/dist/crunch/0.10.0/crunch-0.10.0-cdh5.1.5.tar.gz	https://archive.apache.org/dist/crunch/0.10.0/crunch-0.10.0-cdh5.1.5.html	https://archive.apache.org/dist/crunch/0.10.0/crunch-0.10.0-cdh5.1.5Changes.txt
		Tarball	Release notes	Changes
DataFu	pig-udfdatafu-1.10.0+cdh5.1.5+17	https://archive.apache.org/dist/pig/udfdatafu/1.10.0/cdh5.1.5/pig-udfdatafu-1.10.0-cdh5.1.5.tar.gz	https://archive.apache.org/dist/pig/udfdatafu/1.10.0/cdh5.1.5/pig-udfdatafu-1.10.0-cdh5.1.5.html	https://archive.apache.org/dist/pig/udfdatafu/1.10.0/cdh5.1.5/pig-udfdatafu-1.10.0-cdh5.1.5Changes.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.1.5+21	https://archive.apache.org/dist/flume-ng/1.5.0/cdh5.1.5/flume-ng-1.5.0-cdh5.1.5.tar.gz	https://archive.apache.org/dist/flume-ng/1.5.0/cdh5.1.5/flume-ng-1.5.0-cdh5.1.5.html	https://archive.apache.org/dist/flume-ng/1.5.0/cdh5.1.5/flume-ng-1.5.0-cdh5.1.5Changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.1.5+854	https://archive.apache.org/dist/hadoop/core/hadoop-2.3.0-cdh5.1.5.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-2.3.0-cdh5.1.5.html	https://archive.apache.org/dist/hadoop/core/hadoop-2.3.0-cdh5.1.5Changes.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.1+cdh5.1.5+94	https://archive.apache.org/dist/hbase/0.98.1/cdh5.1.5/hbase-0.98.1-cdh5.1.5.tar.gz	https://archive.apache.org/dist/hbase/0.98.1/cdh5.1.5/hbase-0.98.1-cdh5.1.5.html	https://archive.apache.org/dist/hbase/0.98.1/cdh5.1.5/hbase-0.98.1-cdh5.1.5Changes.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.1.5+18	https://archive.apache.org/dist/hbase/solr/1.5/cdh5.1.5/hbase-solr-1.5-cdh5.1.5.tar.gz	https://archive.apache.org/dist/hbase/solr/1.5/cdh5.1.5/hbase-solr-1.5-cdh5.1.5.html	https://archive.apache.org/dist/hbase/solr/1.5/cdh5.1.5/hbase-solr-1.5-cdh5.1.5Changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.1.5+413	https://archive.apache.org/dist/hive/0.12.0/cdh5.1.5/hive-0.12.0-cdh5.1.5.tar.gz	https://archive.apache.org/dist/hive/0.12.0/cdh5.1.5/hive-0.12.0-cdh5.1.5.html	https://archive.apache.org/dist/hive/0.12.0/cdh5.1.5/hive-0.12.0-cdh5.1.5Changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Parquet-format	parquet-format-1.0.0+cdh5.15+12	https://s3.amazonaws.com/parquet-format-1.0.0+cdh5.15.tar.gz	https://s3.amazonaws.com/parquet-format-1.0.0+cdh5.15.release.html	https://s3.amazonaws.com/parquet-format-1.0.0+cdh5.15.changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.1.5+40	https://s3.amazonaws.com/pig-0.12.0+cdh5.1.5.tar.gz	https://s3.amazonaws.com/pig-0.12.0+cdh5.1.5.release.html	https://s3.amazonaws.com/pig-0.12.0+cdh5.1.5.changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.1.5+0	https://s3.amazonaws.com/search-1.0.0+cdh5.1.5.tar.gz	https://s3.amazonaws.com/search-1.0.0+cdh5.1.5.release.html	https://s3.amazonaws.com/search-1.0.0+cdh5.1.5.changes.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.3.0+cdh5.1.5+168	https://s3.amazonaws.com/sentry-1.3.0+cdh5.1.5.tar.gz	https://s3.amazonaws.com/sentry-1.3.0+cdh5.1.5.release.html	https://s3.amazonaws.com/sentry-1.3.0+cdh5.1.5.changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.1.5+241	https://s3.amazonaws.com/solr-4.4.0+cdh5.1.5.tar.gz	https://s3.amazonaws.com/solr-4.4.0+cdh5.1.5.release.html	https://s3.amazonaws.com/solr-4.4.0+cdh5.1.5.changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.0.0+cdh5.1.5+49	https://s3.amazonaws.com/spark-1.0.0+cdh5.1.5.tar.gz	https://s3.amazonaws.com/spark-1.0.0+cdh5.1.5.release.html	https://s3.amazonaws.com/spark-1.0.0+cdh5.1.5.changes.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.1.5+61	https://s3.amazonaws.com/sqoop-1.4.4+cdh5.1.5.tar.gz	https://s3.amazonaws.com/sqoop-1.4.4+cdh5.1.5.release.html	https://s3.amazonaws.com/sqoop-1.4.4+cdh5.1.5.changes.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.1.5+32	https://s3.amazonaws.com/sqoop2-1.99.3+cdh5.1.5.tar.gz	https://s3.amazonaws.com/sqoop2-1.99.3+cdh5.1.5.release.html	https://s3.amazonaws.com/sqoop2-1.99.3+cdh5.1.5.changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.1.5+15	https://archive.apache.org/dist/cdh/5/whirr-0.9.0+cdh5.1.5.tar.gz	https://archive.apache.org/dist/cdh/5/whirr-0.9.0+cdh5.1.5.html	https://archive.apache.org/dist/cdh/5/whirr-0.9.0+cdh5.1.5.html#changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.1.5+35	https://archive.apache.org/dist/cdh/5/zookeeper-3.4.5+cdh5.1.5.tar.gz	https://archive.apache.org/dist/cdh/5/zookeeper-3.4.5+cdh5.1.5.html	https://archive.apache.org/dist/cdh/5/zookeeper-3.4.5+cdh5.1.5.html#changes
		Tarball	Release notes	Changes

CDH 5.1.4 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.1.4+35	https://archive.apache.org/dist/cdh/5/avro-1.7.5+cdh5.1.4.tar.gz	https://archive.apache.org/dist/cdh/5/avro-1.7.5+cdh5.1.4.html	https://archive.apache.org/dist/cdh/5/avro-1.7.5+cdh5.1.4.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.10.0+cdh5.1.4+19	https://archive.apache.org/dist/cdh/5/crunch-0.10.0+cdh5.1.4.tar.gz	https://archive.apache.org/dist/cdh/5/crunch-0.10.0+cdh5.1.4.html	https://archive.apache.org/dist/cdh/5/crunch-0.10.0+cdh5.1.4.html#changes
		Tarball	Release notes	Changes
DataFu	pig-udf-datafu-1.10+cdh5.1.4+15	https://archive.apache.org/dist/cdh/5/datafu-1.10+cdh5.1.4.tar.gz	https://archive.apache.org/dist/cdh/5/datafu-1.10+cdh5.1.4.html	https://archive.apache.org/dist/cdh/5/datafu-1.10+cdh5.1.4.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.1.4+20	https://archive.apache.org/dist/cdh/5/flume-ng-1.5.0+cdh5.1.4.tar.gz	https://archive.apache.org/dist/cdh/5/flume-ng-1.5.0+cdh5.1.4.html	https://archive.apache.org/dist/cdh/5/flume-ng-1.5.0+cdh5.1.4.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.1.4+842	https://archive.apache.org/dist/cdh/5/hadoop-2.3.0+cdh5.1.4.tar.gz	https://archive.apache.org/dist/cdh/5/hadoop-2.3.0+cdh5.1.4.html	https://archive.apache.org/dist/cdh/5/hadoop-2.3.0+cdh5.1.4.html#changes
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-0.98.1+cdh5.1.4+93	https://www.apache.org/dist/hbase/p/cdh5/cdh/5/hbase-0.98.1+cdh5.1.4.tar.gz	https://www.apache.org/dist/hbase/p/cdh5/cdh/5/hbase-0.98.1+cdh5.1.4-release.html	https://www.apache.org/dist/hbase/p/cdh5/cdh/5/hbase-0.98.1+cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.1.4+17	https://www.apache.org/dist/hbase/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.1.4.tar.gz	https://www.apache.org/dist/hbase/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.1.4-release.html	https://www.apache.org/dist/hbase/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.1.4+399	https://www.apache.org/dist/hive/p/cdh5/cdh/5/hive-0.12.0+cdh5.1.4.tar.gz	https://www.apache.org/dist/hive/p/cdh5/cdh/5/hive-0.12.0+cdh5.1.4-release.html	https://www.apache.org/dist/hive/p/cdh5/cdh/5/hive-0.12.0+cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.1.4+110	https://www.apache.org/dist/hue/p/cdh5/cdh/5/hue-3.6.0+cdh5.1.4.tar.gz	https://www.apache.org/dist/hue/p/cdh5/cdh/5/hue-3.6.0+cdh5.1.4-release.html	https://www.apache.org/dist/hue/p/cdh5/cdh/5/hue-3.6.0+cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.4+cdh5.1.4+0	(none)	https://www.apache.org/dist/impala/p/cdh5/cdh/5/impala-1.4+cdh5.1.4-release.html	https://www.apache.org/dist/impala/p/cdh5/cdh/5/impala-1.4+cdh5.1.4-CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.10.0+cdh5.1.4+125	https://www.apache.org/dist/kite/p/cdh5/cdh/5/kite-0.10.0+cdh5.1.4.tar.gz	https://www.apache.org/dist/kite/p/cdh5/cdh/5/kite-0.10.0+cdh5.1.4-release.html	https://www.apache.org/dist/kite/p/cdh5/cdh/5/kite-0.10.0+cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.1.4+0	https://www.apache.org/dist/llama/p/cdh5/cdh/5/llama-1.0.0+cdh5.1.4.tar.gz	https://www.apache.org/dist/llama/p/cdh5/cdh/5/llama-1.0.0+cdh5.1.4-release.html	https://www.apache.org/dist/llama/p/cdh5/cdh/5/llama-1.0.0+cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.1.4+16	https://www.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.1.4.tar.gz	https://www.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.1.4-release.html	https://www.apache.org/dist/mahout/p/cdh5/cdh/5/mahout-0.9+cdh5.1.4-CHANGES.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.1.4+260	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.1.4.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.1.4.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.1.4CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.2.5+cdh5.1.4+139	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.2.5+cdh5.1.4.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.2.5+cdh5.1.4.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.2.5+cdh5.1.4CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.0.0+cdh5.1.4+11	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-1.0.0+cdh5.1.4.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-1.0.0+cdh5.1.4.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-1.0.0+cdh5.1.4CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.1.4+38	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.1.4.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.1.4.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.1.4CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.1.4+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.1.4.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.1.4.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.1.4CHANGES.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.3.0+cdh5.1.4+166	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.3.0+cdh5.1.4.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.3.0+cdh5.1.4.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.3.0+cdh5.1.4CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.1.4+236	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.1.4.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.1.4.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.1.4CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.0.0+cdh5.1.4+47	https://archive.apache.org/dist/spark/p/cdh5/cdh/	https://archive.apache.org/dist/spark/p/cdh5/cdh/	https://archive.apache.org/dist/spark/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ parquet-1.0.0-cdh5.1.4.tar.gz	5/ parquet-1.0.0-cdh5.1.4-release.html	5/ parquet-1.0.0-cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.1.4+60	https://archive.apache.org/dist/sqoop/sqoop-1.4.4-cdh5.1.4.tar.gz	https://archive.apache.org/dist/sqoop/sqoop-1.4.4-cdh5.1.4-release.html	https://archive.apache.org/dist/sqoop/sqoop-1.4.4-cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.1.4+31	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.3-cdh5.1.4.tar.gz	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.3-cdh5.1.4-release.html	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.3-cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.1.4+14	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.1.4.tar.gz	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.1.4-release.html	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.1.4+34	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.1.4.tar.gz	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.1.4-release.html	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.1.4-CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.1.3 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.1.3+34	https://archive.apache.org/dist/avro/avro-1.7.5-cdh5.1.3.tar.gz	https://archive.apache.org/dist/avro/avro-1.7.5-cdh5.1.3-release.html	https://archive.apache.org/dist/avro/avro-1.7.5-cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.10.0+cdh5.1.3+18	https://archive.apache.org/dist/crunch/crunch-0.10.0-cdh5.1.3.tar.gz	https://archive.apache.org/dist/crunch/crunch-0.10.0-cdh5.1.3-release.html	https://archive.apache.org/dist/crunch/crunch-0.10.0-cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.1.3+13	https://archive.apache.org/dist/pigudatafu/pigudatafu-1.10-cdh5.1.3.tar.gz	https://archive.apache.org/dist/pigudatafu/pigudatafu-1.10-cdh5.1.3-release.html	https://archive.apache.org/dist/pigudatafu/pigudatafu-1.10-cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/cdh5.1.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/cdh5.1.0+cdh5.1.3.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/cdh5.1.0+cdh5.1.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.1.3+15	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.1.3.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.1.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.1.3+824	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.1.3.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.1.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.1+cdh5.1.3+82	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase-0.98.1+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase-0.98.1+cdh5.1.3.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase-0.98.1+cdh5.1.3.CHANGES.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.1.3+16	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.1.3.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hbase-solr-1.5+cdh5.1.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.1.3+381	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive-0.12.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive-0.12.0+cdh5.1.3.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hive-0.12.0+cdh5.1.3.CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.1.3+108	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hue-3.6.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hue-3.6.0+cdh5.1.3.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/hue-3.6.0+cdh5.1.3.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.4+cdh5.1.3+0	(none)	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/impala-1.4+cdh5.1.3.html	https://archive.apache.org/dist/parcels/parcels/p/cdh5/cdh/5/impala-1.4+cdh5.1.3.CHANGES.txt
			Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Kite SDK	kite-0.10.0+cdh5.1.3+124	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/kite-0.10.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/kite-0.10.0+cdh5.1.3.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/kite-0.10.0+cdh5.1.3CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.1.3+0	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/llama-1.0.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/llama-1.0.0+cdh5.1.3.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/llama-1.0.0+cdh5.1.3CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.1.3+15	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/mahout-0.9+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/mahout-0.9+cdh5.1.3.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/mahout-0.9+cdh5.1.3CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.1.3+255	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/oozie-4.0.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/oozie-4.0.0+cdh5.1.3.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/oozie-4.0.0+cdh5.1.3CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.2.5+cdh5.1.3+138	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.2.5+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.2.5+cdh5.1.3.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.2.5+cdh5.1.3CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.0.0+cdh5.1.3+10	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-1.0.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-1.0.0+cdh5.1.3.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-format-1.0.0+cdh5.1.3CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.1.3+37	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/pig-0.12.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/pig-0.12.0+cdh5.1.3.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/pig-0.12.0+cdh5.1.3CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.1.3+0	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/search-1.0.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/search-1.0.0+cdh5.1.3.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/search-1.0.0+cdh5.1.3CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.3.0+cdh5.1.3+162	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.3.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.3.0+cdh5.1.3-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.3.0+cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.1.3+235	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.1.3-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.0.0+cdh5.1.3+45	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.0.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.0.0+cdh5.1.3-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-1.0.0+cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.1.3+59	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.4+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.4+cdh5.1.3-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.4+cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.1.3+30	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.3+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.3+cdh5.1.3-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.3+cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.1.3+13	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.1.3-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.1.3+33	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.1.3.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.1.3-releasenotes.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.1.3-CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.1.2 Packaging and Tarballs



Note:

There is no CDH 5.1.1 release. This skip in the CDH 5.x sequence allows the CDH and Cloudera Manager components of Cloudera Enterprise 5.1.2 to have consistent numbering.

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.1.2+32	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.5-cdh5.1.2.tar.gz	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.5-cdh5.1.2-releasenotes.html	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.5-cdh5.1.2-changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.10.0+cdh5.1.2+16	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.10.0-cdh5.1.2.tar.gz	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.10.0-cdh5.1.2-releasenotes.html	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.10.0-cdh5.1.2-changes.txt
		Tarball	Release notes	Changes
DataFu	pig-udf-datafu-1.10+cdh5.1.2+11	https://archive.apache.org/dist/pig/udf-datafu/p/cdh5/cdh/5/datafu-1.10-cdh5.1.2.tar.gz	https://archive.apache.org/dist/pig/udf-datafu/p/cdh5/cdh/5/datafu-1.10-cdh5.1.2-releasenotes.html	https://archive.apache.org/dist/pig/udf-datafu/p/cdh5/cdh/5/datafu-1.10-cdh5.1.2-changes.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.1.2+13	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.5.0-cdh5.1.2.tar.gz	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.5.0-cdh5.1.2-releasenotes.html	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.5.0-cdh5.1.2-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.1.2+816	https://archive.apache.org/dist/hadoop/core/p/cdh5/cdh/5/hadoop-2.3.0-cdh5.1.2.tar.gz	https://archive.apache.org/dist/hadoop/core/p/cdh5/cdh/5/hadoop-2.3.0-cdh5.1.2-releasenotes.html	https://archive.apache.org/dist/hadoop/core/p/cdh5/cdh/5/hadoop-2.3.0-cdh5.1.2-changes.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.1+cdh5.1.2+70	https://archive.apache.org/dist/hbase/p/cdh5/cdh/5/hbase-0.98.1-cdh5.1.2.tar.gz	https://archive.apache.org/dist/hbase/p/cdh5/cdh/5/hbase-0.98.1-cdh5.1.2-releasenotes.html	https://archive.apache.org/dist/hbase/p/cdh5/cdh/5/hbase-0.98.1-cdh5.1.2-changes.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.1.2+14	https://archive.apache.org/dist/hbase-solr/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.1.2.tar.gz	https://archive.apache.org/dist/hbase-solr/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.1.2-releasenotes.html	https://archive.apache.org/dist/hbase-solr/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.1.2-changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.1.2+375	https://archive.apache.org/dist/hive/p/cdh5/cdh/5/hive-0.12.0-cdh5.1.2.tar.gz	https://archive.apache.org/dist/hive/p/cdh5/cdh/5/hive-0.12.0-cdh5.1.2-releasenotes.html	https://archive.apache.org/dist/hive/p/cdh5/cdh/5/hive-0.12.0-cdh5.1.2-changes.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.1.2+97	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.6.0+cdh5.1.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.6.0+cdh5.1.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hue-3.6.0+cdh5.1.2.PM5.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.4+cdh5.1.2+0	(none)	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-1.4+cdh5.1.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/impala-1.4+cdh5.1.2.PM5.txt
			Release notes	Changes
Kite SDK	kite-0.10.0+cdh5.1.2+122	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-0.10.0+cdh5.1.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-0.10.0+cdh5.1.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/kite-0.10.0+cdh5.1.2.PM5.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.1.2+0	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.1.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.1.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/llama-1.0.0+cdh5.1.2.PM5.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.1.2+13	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/mahout-0.9+cdh5.1.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/mahout-0.9+cdh5.1.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/mahout-0.9+cdh5.1.2.PM5.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.1.2+252	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/oozie-4.0.0+cdh5.1.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/oozie-4.0.0+cdh5.1.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/oozie-4.0.0+cdh5.1.2.PM5.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.25+cdh5.1.2+136	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-1.25+cdh5.1.2.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-1.25+cdh5.1.2.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/parquet-1.25+cdh5.1.2.PM5.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.00+cdh5.1.2+8	https://us.apache.org/dist/hadoop/p/cdh5/cdh/	https://us.apache.org/dist/hadoop/p/cdh5/cdh/	https://us.apache.org/dist/hadoop/p/cdh5/cdh/

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		5/ https://archive.apache.org/dist/parquet/1.0.0+cdh5.1.2+35/parquet-1.0.0+cdh5.1.2+35.tar.gz	5/ https://archive.apache.org/dist/parquet/1.0.0+cdh5.1.2+35/parquet-1.0.0+cdh5.1.2+35.html	5/ https://archive.apache.org/dist/parquet/1.0.0+cdh5.1.2+35/parquet-1.0.0+cdh5.1.2+35.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.1.2+35	https://archive.apache.org/dist/pig/cdh5/cdh/5/pig-0.12.0+cdh5.1.2+35.tar.gz	https://archive.apache.org/dist/pig/cdh5/cdh/5/pig-0.12.0+cdh5.1.2+35.html	https://archive.apache.org/dist/pig/cdh5/cdh/5/pig-0.12.0+cdh5.1.2+35.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.1.2+0	https://archive.apache.org/dist/pig/cdh5/cdh/5/search-1.0.0+cdh5.1.2+0.tar.gz	https://archive.apache.org/dist/pig/cdh5/cdh/5/search-1.0.0+cdh5.1.2+0.html	https://archive.apache.org/dist/pig/cdh5/cdh/5/search-1.0.0+cdh5.1.2+0.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.3.0+cdh5.1.2+160	https://archive.apache.org/dist/pig/cdh5/cdh/5/sentry-1.3.0+cdh5.1.2+160.tar.gz	https://archive.apache.org/dist/pig/cdh5/cdh/5/sentry-1.3.0+cdh5.1.2+160.html	https://archive.apache.org/dist/pig/cdh5/cdh/5/sentry-1.3.0+cdh5.1.2+160.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.1.2+233	https://archive.apache.org/dist/pig/cdh5/cdh/5/solr-4.4.0+cdh5.1.2+233.tar.gz	https://archive.apache.org/dist/pig/cdh5/cdh/5/solr-4.4.0+cdh5.1.2+233.html	https://archive.apache.org/dist/pig/cdh5/cdh/5/solr-4.4.0+cdh5.1.2+233.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.0.0+cdh5.1.2+43	https://archive.apache.org/dist/pig/cdh5/cdh/5/spark-1.0.0+cdh5.1.2+43.tar.gz	https://archive.apache.org/dist/pig/cdh5/cdh/5/spark-1.0.0+cdh5.1.2+43.html	https://archive.apache.org/dist/pig/cdh5/cdh/5/spark-1.0.0+cdh5.1.2+43.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.1.2+57	https://archive.apache.org/dist/pig/cdh5/cdh/5/sqoop-1.4.4+cdh5.1.2+57.tar.gz	https://archive.apache.org/dist/pig/cdh5/cdh/5/sqoop-1.4.4+cdh5.1.2+57.html	https://archive.apache.org/dist/pig/cdh5/cdh/5/sqoop-1.4.4+cdh5.1.2+57.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.1.2+28	https://archive.apache.org/dist/pig/cdh5/cdh/5/sqp2-1.99.3+cdh5.1.2+28.tar.gz	https://archive.apache.org/dist/pig/cdh5/cdh/5/sqp2-1.99.3+cdh5.1.2+28.html	https://archive.apache.org/dist/pig/cdh5/cdh/5/sqp2-1.99.3+cdh5.1.2+28.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Whirr	whirr-0.9.0+cdh5.1.2+11	https://archive.apache.org/dist/whirr/p/cdh5/cdh/5/whirr-0.9.0+cdh5.1.2.tar.gz	https://archive.apache.org/dist/whirr/p/cdh5/cdh/5/whirr-0.9.0+cdh5.1.2.release.html	https://archive.apache.org/dist/whirr/p/cdh5/cdh/5/whirr-0.9.0+cdh5.1.2.CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.1.2+31	https://archive.apache.org/dist/zookeeper/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.1.2.tar.gz	https://archive.apache.org/dist/zookeeper/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.1.2.release.html	https://archive.apache.org/dist/zookeeper/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.1.2.CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.1.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.1.0+30	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.5+cdh5.1.0.tar.gz	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.5+cdh5.1.0.release.html	https://archive.apache.org/dist/avro/p/cdh5/cdh/5/avro-1.7.5+cdh5.1.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.10.0+cdh5.1.0+14	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.10.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.10.0+cdh5.1.0.release.html	https://archive.apache.org/dist/crunch/p/cdh5/cdh/5/crunch-0.10.0+cdh5.1.0.CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pig-udf-datafu-1.1.0+cdh5.1.0+8	https://archive.apache.org/dist/pig-udf-datafu/p/cdh5/cdh/5/datafu-1.1.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/pig-udf-datafu/p/cdh5/cdh/5/datafu-1.1.0+cdh5.1.0.release.html	https://archive.apache.org/dist/pig-udf-datafu/p/cdh5/cdh/5/datafu-1.1.0+cdh5.1.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.5.0+cdh5.1.0+10	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.1.0.release.html	https://archive.apache.org/dist/flume-ng/p/cdh5/cdh/5/flume-ng-1.5.0+cdh5.1.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.1.0+795	https://archive.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.1.0.release.html	https://archive.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.1.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.98.1+cdh5.1.0+64			

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hbase-0.9.1-cdh5.1.0.tar.gz	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hbase-0.9.1-cdh5.1.0-release.html	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hbase-0.9.1-cdh5.1.0-PMS.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.5+cdh5.1.0+12	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.1.0.tar.gz	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.1.0-release.html	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hbase-solr-1.5-cdh5.1.0-PMS.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.1.0+369	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hive-0.12.0-cdh5.1.0.tar.gz	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hive-0.12.0-cdh5.1.0-release.html	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hive-0.12.0-cdh5.1.0-PMS.txt
		Tarball	Release notes	Changes
Hue	hue-3.6.0+cdh5.1.0+86	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hue-3.6.0-cdh5.1.0.tar.gz	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hue-3.6.0-cdh5.1.0-release.html	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/hue-3.6.0-cdh5.1.0-PMS.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.4+cdh5.1.0+0	(none)	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/impala-1.4-cdh5.1.0-release.html	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/impala-1.4-cdh5.1.0-PMS.txt
			Release notes	Changes
Kite SDK	kite-0.10.0+cdh5.1.0+120	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/kite-0.10.0-cdh5.1.0.tar.gz	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/kite-0.10.0-cdh5.1.0-release.html	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/kite-0.10.0-cdh5.1.0-PMS.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.1.0+0	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/llama-1.0.0-cdh5.1.0.tar.gz	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/llama-1.0.0-cdh5.1.0-release.html	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/llama-1.0.0-cdh5.1.0-PMS.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.9+cdh5.1.0+11	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/mahout-0.9-cdh5.1.0.tar.gz	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/mahout-0.9-cdh5.1.0-release.html	https://us.apache.org/dist/techtopology/p/cdh5/cdh/5/mahout-0.9-cdh5.1.0-PMS.txt
			Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.1.0+249	https://archive.apache.org/dist/oozie/oozie-4.0.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/oozie/oozie-4.0.0+cdh5.1.0.html	https://archive.apache.org/dist/oozie/oozie-4.0.0+cdh5.1.0/changes.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.2.5+cdh5.1.0+130	https://archive.apache.org/dist/parquet/parquet-1.2.5+cdh5.1.0.tar.gz	https://archive.apache.org/dist/parquet/parquet-1.2.5+cdh5.1.0.html	https://archive.apache.org/dist/parquet/parquet-1.2.5+cdh5.1.0/changes.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.0.0+cdh5.1.0+6	https://archive.apache.org/dist/parquet-format/parquet-format-1.0.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/parquet-format/parquet-format-1.0.0+cdh5.1.0.html	https://archive.apache.org/dist/parquet-format/parquet-format-1.0.0+cdh5.1.0/changes.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.1.0+33	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.1.0.html	https://archive.apache.org/dist/pig/pig-0.12.0+cdh5.1.0/changes.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.1.0+0	https://archive.apache.org/dist/search/search-1.0.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/search/search-1.0.0+cdh5.1.0.html	https://archive.apache.org/dist/search/search-1.0.0+cdh5.1.0/changes.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.3.0+cdh5.1.0+155	https://archive.apache.org/dist/sentry/sentry-1.3.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/sentry/sentry-1.3.0+cdh5.1.0.html	https://archive.apache.org/dist/sentry/sentry-1.3.0+cdh5.1.0/changes.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.1.0+231	https://archive.apache.org/dist/solr/solr-4.4.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/solr/solr-4.4.0+cdh5.1.0.html	https://archive.apache.org/dist/solr/solr-4.4.0+cdh5.1.0/changes.txt
		Tarball	Release notes	Changes
Apache Spark	spark-1.0.0+cdh5.1.0+41	https://archive.apache.org/dist/spark/spark-1.0.0+cdh5.1.0.tar.gz	https://archive.apache.org/dist/spark/spark-1.0.0+cdh5.1.0.html	https://archive.apache.org/dist/spark/spark-1.0.0+cdh5.1.0/changes.txt

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		5/parcels-1.0.0-cdh5.1.0.tar.gz Tarball	5/parcels-1.0.0-cdh5.1.0-release.html Release notes	5/parcels-1.0.0-cdh5.1.0-changes.txt Changes
Apache Sqoop	sqoop-1.4.4+cdh5.1.0+55	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.4-cdh5.1.0.tar.gz Tarball	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.4-cdh5.1.0-release.html Release notes	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.4-cdh5.1.0-changes.txt Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.1.0+26	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.3-cdh5.1.0.tar.gz Tarball	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.3-cdh5.1.0-release.html Release notes	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.3-cdh5.1.0-changes.txt Changes
Apache Whirr	whirr-0.9.0+cdh5.1.0+9	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0-cdh5.1.0.tar.gz Tarball	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0-cdh5.1.0-release.html Release notes	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0-cdh5.1.0-changes.txt Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.1.0+29	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5-cdh5.1.0.tar.gz Tarball	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5-cdh5.1.0-release.html Release notes	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5-cdh5.1.0-changes.txt Changes

CDH 5.0.x Packaging and Tarball Information

The package version numbers of the projects comprising each CDH 5.0.x release are listed in the following table. For the same type of information for other CDH releases, see [Packaging and Tarballs](#). To view the overall release notes for CDH 5.x.x, see the [Release Notes](#).



Important: When installing CDH from Cloudera tarballs, note that certain features of Hadoop have not been included in the tarballs and will only be available if you have the native libraries present.

To see the details of all the changes and bug-fixes for a given component in a given release, make sure you read the Changes file as well as the Release Notes, following the links in the tables below.

CDH 5.0.6 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.0.6+27	https://archive.apache.org/dist/avro/1.7.5+cdh5.0.6/avro-1.7.5+cdh5.0.6.tar.gz	https://archive.apache.org/dist/avro/1.7.5+cdh5.0.6/avro-1.7.5+cdh5.0.6-releasenotes.html	https://archive.apache.org/dist/avro/1.7.5+cdh5.0.6/avro-1.7.5+cdh5.0.6-changes.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.9.0+cdh5.0.6+32	https://archive.apache.org/dist/crunch/0.9.0+cdh5.0.6/crunch-0.9.0+cdh5.0.6.tar.gz	https://archive.apache.org/dist/crunch/0.9.0+cdh5.0.6/crunch-0.9.0+cdh5.0.6-releasenotes.html	https://archive.apache.org/dist/crunch/0.9.0+cdh5.0.6/crunch-0.9.0+cdh5.0.6-changes.txt
		Tarball	Release notes	Changes
DataFu	pig-udfdatafu-1.10+cdh5.0.6+29	https://archive.apache.org/dist/pig/1.10+cdh5.0.6/pig-udfdatafu-1.10+cdh5.0.6.tar.gz	https://archive.apache.org/dist/pig/1.10+cdh5.0.6/pig-udfdatafu-1.10+cdh5.0.6-releasenotes.html	https://archive.apache.org/dist/pig/1.10+cdh5.0.6/pig-udfdatafu-1.10+cdh5.0.6-changes.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.4.0+cdh5.0.6+133	https://archive.apache.org/dist/flume-ng/1.4.0+cdh5.0.6/flume-ng-1.4.0+cdh5.0.6.tar.gz	https://archive.apache.org/dist/flume-ng/1.4.0+cdh5.0.6/flume-ng-1.4.0+cdh5.0.6-releasenotes.html	https://archive.apache.org/dist/flume-ng/1.4.0+cdh5.0.6/flume-ng-1.4.0+cdh5.0.6-changes.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.0.6+636	https://archive.apache.org/dist/hadoop-core/2.3.0+cdh5.0.6/hadoop-2.3.0+cdh5.0.6.tar.gz	https://archive.apache.org/dist/hadoop-core/2.3.0+cdh5.0.6/hadoop-2.3.0+cdh5.0.6-releasenotes.html	https://archive.apache.org/dist/hadoop-core/2.3.0+cdh5.0.6/hadoop-2.3.0+cdh5.0.6-changes.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.96.1.1+cdh5.0.6+88	https://archive.apache.org/dist/hbase/0.96.1.1+cdh5.0.6/hbase-0.96.1.1+cdh5.0.6.tar.gz	https://archive.apache.org/dist/hbase/0.96.1.1+cdh5.0.6/hbase-0.96.1.1+cdh5.0.6-releasenotes.html	https://archive.apache.org/dist/hbase/0.96.1.1+cdh5.0.6/hbase-0.96.1.1+cdh5.0.6-changes.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.3+cdh5.0.6+52	https://archive.apache.org/dist/hbase-solr/1.3+cdh5.0.6/hbase-solr-1.3+cdh5.0.6.tar.gz	https://archive.apache.org/dist/hbase-solr/1.3+cdh5.0.6/hbase-solr-1.3+cdh5.0.6-releasenotes.html	https://archive.apache.org/dist/hbase-solr/1.3+cdh5.0.6/hbase-solr-1.3+cdh5.0.6-changes.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.0.6+333	https://archive.apache.org/dist/hive/0.12.0+cdh5.0.6/hive-0.12.0+cdh5.0.6.tar.gz	https://archive.apache.org/dist/hive/0.12.0+cdh5.0.6/hive-0.12.0+cdh5.0.6-releasenotes.html	https://archive.apache.org/dist/hive/0.12.0+cdh5.0.6/hive-0.12.0+cdh5.0.6-changes.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		5/ hue-0.12.0-cdh5.0.6.tar.gz	5/ hue-0.12.0-cdh5.0.6-release.html	5/ hue-0.12.0-cdh5.0.6-CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.5.0+cdh5.0.6+391	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ hue-3.5.0-cdh5.0.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ hue-3.5.0-cdh5.0.6-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ hue-3.5.0-cdh5.0.6-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.3+cdh5.0.6+0	(none)	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ impala-1.3-cdh5.0.6-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ impala-1.3-cdh5.0.6-CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.10.0+cdh5.0.6+93	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ kite-0.10.0-cdh5.0.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ kite-0.10.0-cdh5.0.6-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ kite-0.10.0-cdh5.0.6-CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.0.6+0	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.0.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.0.6-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ llama-1.0.0-cdh5.0.6-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.8+cdh5.0.6+38	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ mahout-0.8-cdh5.0.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ mahout-0.8-cdh5.0.6-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ mahout-0.8-cdh5.0.6-CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.0.6+195	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ oozie-4.0.0-cdh5.0.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ oozie-4.0.0-cdh5.0.6-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ oozie-4.0.0-cdh5.0.6-CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.2.5+cdh5.0.6+129	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ parquet-1.2.5-cdh5.0.6.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ parquet-1.2.5-cdh5.0.6-release.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ parquet-1.2.5-cdh5.0.6-CHANGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Parquet-format	parquet-format-1.0.0+cdh5.0.6+14	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ parquet-format-1.0.0+cdh5.0.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ parquet-format-1.0.0+cdh5.0.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ parquet-format-1.0.0+cdh5.0.6.html
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.0.6+41	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.0.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.0.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ pig-0.12.0+cdh5.0.6.html
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.0.6+0	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.0.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.0.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ search-1.0.0+cdh5.0.6.html
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.2.0+cdh5.0.6+82	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sentry-1.2.0+cdh5.0.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sentry-1.2.0+cdh5.0.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sentry-1.2.0+cdh5.0.6.html
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.0.6+204	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.0.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.0.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ solr-4.4.0+cdh5.0.6.html
		Tarball	Release notes	Changes
Apache Spark	spark-0.9.0+cdh5.0.6+44	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-0.9.0+cdh5.0.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-0.9.0+cdh5.0.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ spark-0.9.0+cdh5.0.6.html
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.0.6+55	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.4+cdh5.0.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.4+cdh5.0.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop-1.4.4+cdh5.0.6.html
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.0.6+37	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.0.6.tar.gz	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.0.6.html	https:// us.apache.org/dist/mirrors/ p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.0.6.html
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.0.6+18	https://archive.apache.org/dist/cdh/5/whirr-0.9.0+cdh5.0.6.tar.gz	https://archive.apache.org/dist/cdh/5/whirr-0.9.0+cdh5.0.6.html	https://archive.apache.org/dist/cdh/5/whirr-0.9.0+cdh5.0.6.PAGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.0.6+40	https://archive.apache.org/dist/cdh/5/zookeeper-3.4.5+cdh5.0.6.tar.gz	https://archive.apache.org/dist/cdh/5/zookeeper-3.4.5+cdh5.0.6.html	https://archive.apache.org/dist/cdh/5/zookeeper-3.4.5+cdh5.0.6.PAGES.txt
		Tarball	Release notes	Changes

CDH 5.0.5 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.0.5+26	https://archive.apache.org/dist/cdh/5/avro-1.7.5+cdh5.0.5.tar.gz	https://archive.apache.org/dist/cdh/5/avro-1.7.5+cdh5.0.5.html	https://archive.apache.org/dist/cdh/5/avro-1.7.5+cdh5.0.5.PAGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.9.0+cdh5.0.5+31	https://archive.apache.org/dist/cdh/5/crunch-0.9.0+cdh5.0.5.tar.gz	https://archive.apache.org/dist/cdh/5/crunch-0.9.0+cdh5.0.5.html	https://archive.apache.org/dist/cdh/5/crunch-0.9.0+cdh5.0.5.PAGES.txt
		Tarball	Release notes	Changes
DataFu	pig-udf-datafu-1.10+cdh5.0.5+27	https://archive.apache.org/dist/cdh/5/datafu-1.10+cdh5.0.5.tar.gz	https://archive.apache.org/dist/cdh/5/datafu-1.10+cdh5.0.5.html	https://archive.apache.org/dist/cdh/5/datafu-1.10+cdh5.0.5.PAGES.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.4.0+cdh5.0.5+132	https://archive.apache.org/dist/cdh/5/flume-ng-1.4.0+cdh5.0.5.tar.gz	https://archive.apache.org/dist/cdh/5/flume-ng-1.4.0+cdh5.0.5.html	https://archive.apache.org/dist/cdh/5/flume-ng-1.4.0+cdh5.0.5.PAGES.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.0.5+626	https://archive.apache.org/dist/cdh/5/hadoop-2.3.0+cdh5.0.5.tar.gz	https://archive.apache.org/dist/cdh/5/hadoop-2.3.0+cdh5.0.5.html	https://archive.apache.org/dist/cdh/5/hadoop-2.3.0+cdh5.0.5.PAGES.txt
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache HBase	hbase-0.96.1.1+cdh5.0.5+87	https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.96.1.1+cdh5.0.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.96.1.1+cdh5.0.5.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.96.1.1+cdh5.0.5.html
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.3+cdh5.0.5+51	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.3+cdh5.0.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.3+cdh5.0.5.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.3+cdh5.0.5.html
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.0.5+330	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.12.0+cdh5.0.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.12.0+cdh5.0.5.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.12.0+cdh5.0.5.html
		Tarball	Release notes	Changes
Hue	hue-3.5.0+cdh5.0.5+390	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0+cdh5.0.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0+cdh5.0.5.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0+cdh5.0.5.html
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.3+cdh5.0.5+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-1.3+cdh5.0.5.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-1.3+cdh5.0.5.html
			Release notes	Changes
Kite SDK	kite-0.10.0+cdh5.0.5+92	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.5.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.5.html
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.0.5+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.5.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.5.html
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.8+cdh5.0.5+37	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.5.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.5.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.5.html

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.0.5+194	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.0.5.tar.gz	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.0.5.html	https://archive.apache.org/dist/oozie/p/cdh5/cdh/5/oozie-4.0.0+cdh5.0.5.PAHS.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.2.5+cdh5.0.5+128	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.2.5+cdh5.0.5.tar.gz	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.2.5+cdh5.0.5.html	https://archive.apache.org/dist/parquet/p/cdh5/cdh/5/parquet-1.2.5+cdh5.0.5.PAHS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.0.0+cdh5.0.5+13	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.5.tar.gz	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.5.html	https://archive.apache.org/dist/parquet-format/p/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.5.PAHS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.0.5+40	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.0.5.tar.gz	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.0.5.html	https://archive.apache.org/dist/pig/p/cdh5/cdh/5/pig-0.12.0+cdh5.0.5.PAHS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.0.5+0	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.0.5.tar.gz	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.0.5.html	https://archive.apache.org/dist/search/p/cdh5/cdh/5/search-1.0.0+cdh5.0.5.PAHS.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.2.0+cdh5.0.5+81	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.2.0+cdh5.0.5.tar.gz	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.2.0+cdh5.0.5.html	https://archive.apache.org/dist/sentry/p/cdh5/cdh/5/sentry-1.2.0+cdh5.0.5.PAHS.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.0.5+200	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.0.5.tar.gz	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.0.5.html	https://archive.apache.org/dist/solr/p/cdh5/cdh/5/solr-4.4.0+cdh5.0.5.PAHS.txt
		Tarball	Release notes	Changes
Apache Spark	spark-0.9.0+cdh5.0.5+43	https://archive.apache.org/dist/spark/p/cdh5/cdh/	https://archive.apache.org/dist/spark/p/cdh5/cdh/	https://archive.apache.org/dist/spark/p/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ sqoop-0.90+cdh5.0.5.tar.gz	5/ sqoop-0.90+cdh5.0.5.html	5/ sqoop-0.90+cdh5.0.5.PM5.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.0.5+54	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ sqoop-1.4.4+cdh5.0.5.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ sqoop-1.4.4+cdh5.0.5.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ sqoop-1.4.4+cdh5.0.5.PM5.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.0.5+36	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.0.5.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.0.5.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.0.5.PM5.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.0.5+17	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.0.5.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.0.5.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ whirr-0.9.0+cdh5.0.5.PM5.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.0.5+39	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.0.5.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.0.5.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ zookeeper-3.4.5+cdh5.0.5.PM5.txt
		Tarball	Release notes	Changes

CDH 5.0.4 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.0.4+23	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ avro-1.7.5+cdh5.0.4.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ avro-1.7.5+cdh5.0.4.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ avro-1.7.5+cdh5.0.4.PM5.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.9.0+cdh5.0.4+28	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ crunch-0.9.0+cdh5.0.4.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ crunch-0.9.0+cdh5.0.4.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ crunch-0.9.0+cdh5.0.4.PM5.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.0.4+24	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ pigudatafu-1.10+cdh5.0.4.tar.gz	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ pigudatafu-1.10+cdh5.0.4.html	https:// s3.amazonaws.com/ p/cdh5/cdh/ 5/ pigudatafu-1.10+cdh5.0.4.PM5.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/cdh-1.10+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/cdh-1.10+cdh5.0.4.html	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/cdh-1.10+cdh5.0.4.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.4.0+cdh5.0.4+121	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/flume-ng-1.4.0+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/flume-ng-1.4.0+cdh5.0.4.html	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/flume-ng-1.4.0+cdh5.0.4.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.0.4+616	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.0.4.html	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.0.4.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-0.96.1.1+cdh5.0.4+81	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hbase-0.96.1.1+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hbase-0.96.1.1+cdh5.0.4.html	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hbase-0.96.1.1+cdh5.0.4.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.3+cdh5.0.4+48	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hbase-solr-1.3+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hbase-solr-1.3+cdh5.0.4.html	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hbase-solr-1.3+cdh5.0.4.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.0.4+325	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hive-0.12.0+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hive-0.12.0+cdh5.0.4.html	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hive-0.12.0+cdh5.0.4.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.5.0+cdh5.0.4+386	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hue-3.5.0+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hue-3.5.0+cdh5.0.4.html	https://archive.apache.org/dist/parcels/parcels/5/p/cdh5/cdh/5/hue-3.5.0+cdh5.0.4.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.3.2+cdh5.0.4+0	(none)	http://archive.apache.org/dist/parcels/parcels/cdh5/cdh/5/impala-1.3.2+cdh5.0.4.html	http://archive.apache.org/dist/parcels/parcels/cdh5/cdh/5/impala-1.3.2+cdh5.0.4.html#changes
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.2.0+cdh5.0.4+78	https://archive.apache.org/dist/parcels/cdh/5/sentry-1.2.0+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/cdh/5/sentry-1.2.0+cdh5.0.4.html	https://archive.apache.org/dist/parcels/cdh/5/sentry-1.2.0+cdh5.0.4CHANGES.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.0.4+197	https://archive.apache.org/dist/parcels/cdh/5/solr-4.4.0+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/cdh/5/solr-4.4.0+cdh5.0.4.html	https://archive.apache.org/dist/parcels/cdh/5/solr-4.4.0+cdh5.0.4CHANGES.txt
		Tarball	Release notes	Changes
Apache Spark	spark-0.9.0+cdh5.0.4+40	https://archive.apache.org/dist/parcels/cdh/5/spark-0.9.0+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/cdh/5/spark-0.9.0+cdh5.0.4.html	https://archive.apache.org/dist/parcels/cdh/5/spark-0.9.0+cdh5.0.4CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.0.4+51	https://archive.apache.org/dist/parcels/cdh/5/sqoop-1.4.4+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/cdh/5/sqoop-1.4.4+cdh5.0.4.html	https://archive.apache.org/dist/parcels/cdh/5/sqoop-1.4.4+cdh5.0.4CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.0.4+33	https://archive.apache.org/dist/parcels/cdh/5/sqoop2-1.99.3+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/cdh/5/sqoop2-1.99.3+cdh5.0.4.html	https://archive.apache.org/dist/parcels/cdh/5/sqoop2-1.99.3+cdh5.0.4CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.0.4+14	https://archive.apache.org/dist/parcels/cdh/5/whirr-0.9.0+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/cdh/5/whirr-0.9.0+cdh5.0.4.html	https://archive.apache.org/dist/parcels/cdh/5/whirr-0.9.0+cdh5.0.4CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.0.4+36	https://archive.apache.org/dist/parcels/cdh/5/zookeeper-3.4.5+cdh5.0.4.tar.gz	https://archive.apache.org/dist/parcels/cdh/5/zookeeper-3.4.5+cdh5.0.4.html	https://archive.apache.org/dist/parcels/cdh/5/zookeeper-3.4.5+cdh5.0.4CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.0.3 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
-----------	-----------------	---------	---------------	--------------

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.0.3+22	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/avro-1.7.5-cdh5.0.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/avro-1.7.5-cdh5.0.3.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/avro-1.7.5-cdh5.0.3.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.9.0+cdh5.0.3+27	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/crunch-0.9.0-cdh5.0.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/crunch-0.9.0-cdh5.0.3.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/crunch-0.9.0-cdh5.0.3.html#changes
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.0.3+20	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/datafu-1.10-cdh5.0.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/datafu-1.10-cdh5.0.3.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/datafu-1.10-cdh5.0.3.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.4.0+cdh5.0.3+120	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-ng-1.4.0-cdh5.0.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-ng-1.4.0-cdh5.0.3.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/flume-ng-1.4.0-cdh5.0.3.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.0.3+597	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.3.0-cdh5.0.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.3.0-cdh5.0.3.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hadoop-2.3.0-cdh5.0.3.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-0.96.1.1+cdh5.0.3+78	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-0.96.1.1-cdh5.0.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-0.96.1.1-cdh5.0.3.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-0.96.1.1-cdh5.0.3.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.3+cdh5.0.3+47	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.3-cdh5.0.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.3-cdh5.0.3.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hbase-solr-1.3-cdh5.0.3.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.0.3+322	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-0.12.0-cdh5.0.3.tar.gz	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-0.12.0-cdh5.0.3.html	https://archive.apache.org/dist/hadoop/core/hadoop-5/p/cdh5/cdh/5/hive-0.12.0-cdh5.0.3.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-3.5.0+cdh5.0.3+381	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0+cdh5.0.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0+cdh5.0.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0+cdh5.0.3.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.3+cdh5.0.3+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-1.3+cdh5.0.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-1.3+cdh5.0.3.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.10.0+cdh5.0.3+88	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.3.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.0.3+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.8+cdh5.0.3+33	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.3.CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.0.3+185	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.3.CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.2.5+cdh5.0.3+97	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.3.CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.0.0+cdh5.0.3+9	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.3.release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.3.CHANGES.txt

Component	Package Version	Tarball	Release Notes	Changes File
		5/ pig-0.12.0+cdh5.0.3.tar.gz	5/ release-notes	5/ changes
Apache Pig	pig-0.12.0+cdh5.0.3+36	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ pig-0.12.0+cdh5.0.3.tar.gz	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ release-notes	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ changes
Cloudera Search	search-1.0.0+cdh5.0.3+0	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ search-1.0.0+cdh5.0.3.tar.gz	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ release-notes	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ changes
Apache Sentry (incubating)	sentry-1.2.0+cdh5.0.3+77	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ sentry-1.2.0+cdh5.0.3.tar.gz	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ release-notes	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ changes
Apache Solr	solr-4.4.0+cdh5.0.3+196	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ solr-4.4.0+cdh5.0.3.tar.gz	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ release-notes	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ changes
Apache Spark	spark-0.9.0+cdh5.0.3+38	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ spark-0.9.0+cdh5.0.3.tar.gz	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ release-notes	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ changes
Apache Sqoop	sqoop-1.4.4+cdh5.0.3+50	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ sqoop-1.4.4+cdh5.0.3.tar.gz	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ release-notes	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.0.3+32	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ sqoop2-1.99.3+cdh5.0.3.tar.gz	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ release-notes	https:// artifacts.elastic.co/ packages/cdh5/cdh/ 5/ changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
Apache Whirr	whirr-0.9.0+cdh5.0.3+13	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.0.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.0.3.html	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.0.3.PM5.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.0.3+35	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.0.3.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.0.3.html	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.0.3.PM5.txt
		Tarball	Release notes	Changes

CDH 5.0.2 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.0.2+20	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.5+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.5+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/avro-1.7.5+cdh5.0.2.PM5.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.9.0+cdh5.0.2+25	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.9.0+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.9.0+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/crunch-0.9.0+cdh5.0.2.PM5.txt
		Tarball	Release notes	Changes
DataFu	pigudatafu-1.10+cdh5.0.2+18	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/datafu-1.10+cdh5.0.2.PM5.txt
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.4.0+cdh5.0.2+115	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.4.0+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.4.0+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/flume-ng-1.4.0+cdh5.0.2.PM5.txt
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.0.2+579	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.3.0+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.3.0+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/hadoop-2.3.0+cdh5.0.2.PM5.txt
		Tarball	Release notes	Changes
Apache HBase	hbase-0.96.1.1+cdh5.0.2+74			

Component	Package Version	Tarball	Release Notes	Changes File
		https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.95.1.1-cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.95.1.1-cdh5.0.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-0.95.1.1-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.3+cdh5.0.2+42	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.3-cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.3-cdh5.0.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hbase-solr-1.3-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.0.2+319	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.12.0-cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.12.0-cdh5.0.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hive-0.12.0-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
Hue	hue-3.5.0+cdh5.0.2+373	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0-cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0-cdh5.0.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.3+cdh5.0.2+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-1.3-cdh5.0.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/impala-1.3-cdh5.0.2-CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.10.0+cdh5.0.2+86	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0-cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0-cdh5.0.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.0.2+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.0.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.8+cdh5.0.2+31	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8-cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8-cdh5.0.2-release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8-cdh5.0.2-CHANGES.txt
			Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.0.2+180	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.2.PMS.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.2.5+cdh5.0.2+95	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.2.PMS.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.0.0+cdh5.0.2+7	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.2.PMS.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.0.2+34	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.0.2.PMS.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.0.2+0	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.0.2.PMS.txt
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.2.0+cdh5.0.2+75	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.2.0+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.2.0+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/sentry-1.2.0+cdh5.0.2.PMS.txt
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.0.2+190	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.4.0+cdh5.0.2.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.4.0+cdh5.0.2.html	https://s3.amazonaws.com/cdh5/cdh/5/solr-4.4.0+cdh5.0.2.PMS.txt
		Tarball	Release notes	Changes
Apache Spark	spark-0.9.0+cdh5.0.2+36	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ parquet-0.9.0-cdh5.0.2.tar.gz	5/ parquet-0.9.0-cdh5.0.2-release.html	5/ parquet-0.9.0-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.0.2+48	https://archive.apache.org/dist/sqoop/sqoop-1.4.4-cdh5.0.2.tar.gz	https://archive.apache.org/dist/sqoop/sqoop-1.4.4-cdh5.0.2-release.html	https://archive.apache.org/dist/sqoop/sqoop-1.4.4-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.0.2+30	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.3-cdh5.0.2.tar.gz	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.3-cdh5.0.2-release.html	https://archive.apache.org/dist/sqoop2/sqoop2-1.99.3-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.0.2+11	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.0.2.tar.gz	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.0.2-release.html	https://archive.apache.org/dist/whirr/whirr-0.9.0-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.0.2+33	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.0.2.tar.gz	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.0.2-release.html	https://archive.apache.org/dist/zookeeper/zookeeper-3.4.5-cdh5.0.2-CHANGES.txt
		Tarball	Release notes	Changes

CDH 5.0.1 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.0.1+17	https://archive.apache.org/dist/avro/avro-1.7.5-cdh5.0.1.tar.gz	https://archive.apache.org/dist/avro/avro-1.7.5-cdh5.0.1-release.html	https://archive.apache.org/dist/avro/avro-1.7.5-cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.9.0+cdh5.0.1+22	https://archive.apache.org/dist/crunch/crunch-0.9.0-cdh5.0.1.tar.gz	https://archive.apache.org/dist/crunch/crunch-0.9.0-cdh5.0.1-release.html	https://archive.apache.org/dist/crunch/crunch-0.9.0-cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes
DataFu	pig-udf-datafu-1.10+cdh5.0.1+8	https://archive.apache.org/dist/pig-udf-datafu/pig-udf-datafu-1.10-cdh5.0.1.tar.gz	https://archive.apache.org/dist/pig-udf-datafu/pig-udf-datafu-1.10-cdh5.0.1-release.html	https://archive.apache.org/dist/pig-udf-datafu/pig-udf-datafu-1.10-cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		https://www.cloudera.com/p/cdh5/cdh/5/cdh5.0.110+cdh5.0.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/cdh5.0.110+cdh5.0.1.html	https://www.cloudera.com/p/cdh5/cdh/5/cdh5.0.110+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.4.0+cdh5.0.1+112	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.4.0+cdh5.0.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.4.0+cdh5.0.1.html	https://www.cloudera.com/p/cdh5/cdh/5/flume-ng-1.4.0+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.0.1+566	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.0.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.0.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hadoop-2.3.0+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-0.96.1.1+cdh5.0.1+67	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.96.1.1+cdh5.0.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.96.1.1+cdh5.0.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-0.96.1.1+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.3+cdh5.0.1+39	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.3+cdh5.0.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.3+cdh5.0.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hbase-solr-1.3+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.0.1+312	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.12.0+cdh5.0.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.12.0+cdh5.0.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hive-0.12.0+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Hue	hue-3.5.0+cdh5.0.1+370	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.5.0+cdh5.0.1.tar.gz	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.5.0+cdh5.0.1.html	https://www.cloudera.com/p/cdh5/cdh/5/hue-3.5.0+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.3+cdh5.0.1+0	(none)	https://www.cloudera.com/p/cdh5/cdh/5/impala-1.3+cdh5.0.1.html	https://www.cloudera.com/p/cdh5/cdh/5/impala-1.3+cdh5.0.1.html#changes
			Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Kite SDK	kite-0.10.0+cdh5.0.1+80	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.0.1+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.8+cdh5.0.1+28	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.0.1+177	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.2.5+cdh5.0.1+92	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.0.0+cdh5.0.1+4	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/parquet-format-1.0.0+cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.0.1+30	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.0.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.0.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/pig-0.12.0+cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.0.1+0	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.0.1.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.0.1-release.html	https://s3.amazonaws.com/cdh5/cdh/5/search-1.0.0+cdh5.0.1-CHANGES.txt
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.2.0+cdh5.0.1+72	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.2.0+cdh5.0.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.2.0+cdh5.0.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sentry-1.2.0+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.0.1+187	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.0.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.0.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/solr-4.4.0+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache Spark	spark-0.9.0+cdh5.0.1+32	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-0.9.0+cdh5.0.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-0.9.0+cdh5.0.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/spark-0.9.0+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.0.1+44	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.4+cdh5.0.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.4+cdh5.0.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop-1.4.4+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.0.1+27	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.3+cdh5.0.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.3+cdh5.0.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/sqoop2-1.99.3+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache Whirr	whirr-0.9.0+cdh5.0.1+7	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.0.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.0.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/whirr-0.9.0+cdh5.0.1.html#changes
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.0.1+29	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.0.1.tar.gz	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.0.1.html	https://archive.apache.org/dist/parcels/p/cdh5/cdh/5/zookeeper-3.4.5+cdh5.0.1.html#changes
		Tarball	Release notes	Changes

CDH 5.0.0 Packaging and Tarballs

Component	Package Version	Tarball	Release Notes	Changes File
-----------	-----------------	---------	---------------	--------------

Component	Package Version	Tarball	Release Notes	Changes File
Apache Avro	avro-1.7.5+cdh5.0.0+16	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/avro-1.7.5-cdh5.0.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/avro-1.7.5-cdh5.0.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/avro-1.7.5-cdh5.0.0.html#changes
		Tarball	Release notes	Changes
Apache Crunch	crunch-0.9.0+cdh5.0.0+21	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/crunch-0.9.0-cdh5.0.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/crunch-0.9.0-cdh5.0.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/crunch-0.9.0-cdh5.0.0.html#changes
		Tarball	Release notes	Changes
DataFu	pig-udf-datafu-1.1.0+cdh5.0.0+7	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/datafu-1.1.0-cdh5.0.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/datafu-1.1.0-cdh5.0.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/datafu-1.1.0-cdh5.0.0.html#changes
		Tarball	Release notes	Changes
Apache Flume	flume-ng-1.4.0+cdh5.0.0+111	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/flume-ng-1.4.0-cdh5.0.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/flume-ng-1.4.0-cdh5.0.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/flume-ng-1.4.0-cdh5.0.0.html#changes
		Tarball	Release notes	Changes
Apache Hadoop	hadoop-2.3.0+cdh5.0.0+548	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.3.0-cdh5.0.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.3.0-cdh5.0.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hadoop-2.3.0-cdh5.0.0.html#changes
		Tarball	Release notes	Changes
Apache HBase	hbase-0.96.1.1+cdh5.0.0+60	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-0.96.1.1-cdh5.0.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-0.96.1.1-cdh5.0.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-0.96.1.1-cdh5.0.0.html#changes
		Tarball	Release notes	Changes
HBase-Solr	hbase-solr-1.3+cdh5.0.0+38	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-solr-1.3-cdh5.0.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-solr-1.3-cdh5.0.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hbase-solr-1.3-cdh5.0.0.html#changes
		Tarball	Release notes	Changes
Apache Hive	hive-0.12.0+cdh5.0.0+308	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-0.12.0-cdh5.0.0.tar.gz	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-0.12.0-cdh5.0.0.html	https://us.apache.org/dist/hadoop/p/cdh5/cdh/5/hive-0.12.0-cdh5.0.0.html#changes
		Tarball	Release notes	Changes

Version and Download Information

Component	Package Version	Tarball	Release Notes	Changes File
		Tarball	Release notes	Changes
Hue	hue-3.5.0+cdh5.0.0+365	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0+cdh5.0.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0+cdh5.0.0.rst	https://s3.amazonaws.com/cdh5/cdh/5/hue-3.5.0+cdh5.0.0.CHANGES.txt
		Tarball	Release notes	Changes
Cloudera Impala	impala-1.3.0+cdh5.0.0+0	(none)	https://s3.amazonaws.com/cdh5/cdh/5/impala-1.3.0+cdh5.0.0.rst	https://s3.amazonaws.com/cdh5/cdh/5/impala-1.3.0+cdh5.0.0.CHANGES.txt
			Release notes	Changes
Kite SDK	kite-0.10.0+cdh5.0.0+79	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.0.rst	https://s3.amazonaws.com/cdh5/cdh/5/kite-0.10.0+cdh5.0.0.CHANGES.txt
		Tarball	Release notes	Changes
Llama	llama-1.0.0+cdh5.0.0+0	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.0.rst	https://s3.amazonaws.com/cdh5/cdh/5/llama-1.0.0+cdh5.0.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Mahout	mahout-0.8+cdh5.0.0+27	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.0.rst	https://s3.amazonaws.com/cdh5/cdh/5/mahout-0.8+cdh5.0.0.CHANGES.txt
		Tarball	Release notes	Changes
Apache Oozie	oozie-4.0.0+cdh5.0.0+174	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.0.rst	https://s3.amazonaws.com/cdh5/cdh/5/oozie-4.0.0+cdh5.0.0.CHANGES.txt
		Tarball	Release notes	Changes
Parquet	parquet-1.2.5+cdh5.0.0+91	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.0.rst	https://s3.amazonaws.com/cdh5/cdh/5/parquet-1.2.5+cdh5.0.0.CHANGES.txt
		Tarball	Release notes	Changes
Parquet-format	parquet-format-1.0.0+cdh5.0.0+3	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/	https://s3.amazonaws.com/cdh5/cdh/

Component	Package Version	Tarball	Release Notes	Changes File
		5/ parquet-1.0.0-cdh5.0.0.tar.gz	5/ parquet-1.0.0-cdh5.0.0-releasenotes.html	5/ parquet-1.0.0-cdh5.0.0-changes.html
		Tarball	Release notes	Changes
Apache Pig	pig-0.12.0+cdh5.0.0+27	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0-cdh5.0.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0-cdh5.0.0-releasenotes.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ pig-0.12.0-cdh5.0.0-changes.html
		Tarball	Release notes	Changes
Cloudera Search	search-1.0.0+cdh5.0.0+0	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0-cdh5.0.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0-cdh5.0.0-releasenotes.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ search-1.0.0-cdh5.0.0-changes.html
		Tarball	Release notes	Changes
Apache Sentry (incubating)	sentry-1.2.0+cdh5.0.0+71	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.2.0-cdh5.0.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.2.0-cdh5.0.0-releasenotes.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sentry-1.2.0-cdh5.0.0-changes.html
		Tarball	Release notes	Changes
Apache Solr	solr-4.4.0+cdh5.0.0+178	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0-cdh5.0.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0-cdh5.0.0-releasenotes.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ solr-4.4.0-cdh5.0.0-changes.html
		Tarball	Release notes	Changes
Apache Spark	spark-0.9.0+cdh5.0.0+31	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-0.9.0-cdh5.0.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-0.9.0-cdh5.0.0-releasenotes.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ spark-0.9.0-cdh5.0.0-changes.html
		Tarball	Release notes	Changes
Apache Sqoop	sqoop-1.4.4+cdh5.0.0+43	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.4-cdh5.0.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.4-cdh5.0.0-releasenotes.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop-1.4.4-cdh5.0.0-changes.html
		Tarball	Release notes	Changes
Apache Sqoop2	sqoop2-1.99.3+cdh5.0.0+26	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.3-cdh5.0.0.tar.gz	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.3-cdh5.0.0-releasenotes.html	https:// archive.apache.org/dist/ p/cdh5/cdh/ 5/ sqoop2-1.99.3-cdh5.0.0-changes.html
		Tarball	Release notes	Changes

Component	Package Version	Tarball	Release Notes	Changes File
Apache Whirr	whirr-0.9.0+cdh5.0.0+4	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.0.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.0.0.html	https://s3.amazonaws.com/cdh5/cdh/5/whirr-0.9.0+cdh5.0.0.txt
		Tarball	Release notes	Changes
Apache ZooKeeper	zookeeper-3.4.5+cdh5.0.0+28	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.0.0.tar.gz	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.0.0.html	https://s3.amazonaws.com/cdh5/cdh/5/zookeeper-3.4.5+cdh5.0.0.txt
		Tarball	Release notes	Changes

Using the CDH 5 Maven Repository

If you want to build applications or tools for use with CDH 5 components and you are using Maven or Ivy for dependency management, you can pull the CDH 5 artifacts from the Cloudera Maven repository. The repository is available at <https://repository.cloudera.com/artifactory/cloudera-repos/>.



Important: When you build an application JAR, *do not* include CDH JARs, because they are already provided. If you do, upgrading CDH can break your application. To avoid this situation, set the Maven dependency scope to `provided`. If you have already built applications which include the CDH JARs, update the dependency to set scope to `provided` and recompile.



Important: If your cluster hosts are running CDH 5.16.x and the Ubuntu 16.04 or Debian 9 operating system and you are using the OpenJDK SDK to compile applications, note the following:

The following OpenJDK versions will cause the Maven Surefire plug-in to fail:

- Ubuntu 16.04: 8u181-b13-1ubuntu0.16.04.1
- Debian 9: 8u181-b13-2~deb9u1

with the following error message:

```
Error: Could not find or load main class
org.apache.maven.surefire.booter.ForkedBooter
```

Cloudera recommends that you use an earlier version of OpenJDK to compile your applications.

For more information, see the following bug reports:

- Debian: <https://bugs.debian.org/cgi-bin/bugreport.cgi?bug=911925>
- Ubuntu: <https://bugs.launchpad.net/ubuntu/+source/openjdk-8/+bug/1800792>

The following is a sample POM (`pom.xml`) file:

```
<project xmlns="http://maven.apache.org/POM/4.0.0"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://maven.apache.org/POM/4.0.0
http://maven.apache.org/maven-v4_0_0.xsd">
<repositories>
<repository>
<id>cloudera</id>
<url>https://repository.cloudera.com/artifactory/cloudera-repos/</url>
</repository>
```

```
</repositories>
</project>
```

Managing Hadoop API Dependencies in CDH 5

In CDH 3, all of the Hadoop API implementations were confined to a single JAR file (`hadoop-core`) plus a few of its dependencies. It was relatively straightforward to make sure that classes from these JAR files were available at runtime.

CDH 4 and CDH 5 are more complex: they bundle both MRv1 and MRv2 (YARN). To simplify things, CDH 4 and CDH 5 provide a Maven-based way of managing client-side Hadoop API dependencies that saves you from having to figure out the exact names and locations of all the JAR files needed to provide Hadoop APIs.

In CDH 5, Cloudera recommends that you use a `hadoop-client` artifact for all clients, instead of managing JAR-file-based dependencies manually.



Important: When you build an application JAR, *do not* include CDH JARs, because they are already provided. If you do, upgrading CDH can break your application. To avoid this situation, set the Maven dependency `scope` to `provided`. If you have already built applications which include the CDH JARs, update the dependency to set `scope` to `provided` and recompile.

Flavors of the `hadoop-client` Artifact

There are two different flavors of the `hadoop-client` artifact: a Maven-based Project Object Model (POM) artifact and a Linux package, `hadoop-client`. The former lets you manage Hadoop API dependencies at both compile and run time for your Maven- or Ivy-based projects; the latter provides a familiar interface in the form of a collection of JAR files that can be added to your classpath directly.

Versions of the `hadoop-client` Artifact

CDH 5 provides two distinct versions of the `hadoop-client` artifact: one for MRv1 and one for MRv2 (YARN). If you're using the Maven-based POM `hadoop-client` artifact, you can use the version string to distinguish between them:

Table 42: CDH Version Strings for Maven

CDH Version	MRv1 Version String	YARN Version String
5.0.0 - 5.1.x	2.3.0-mr1-cdh5.x.x	2.3.0-cdh5.x.x
5.2.0 - 5.3.x	2.5.0-mr1-cdh5.x.x	2.5.0-cdh5.x.x
5.4.0 and higher	2.6.0-mr1-cdh5.x.x	2.6.0-cdh5.x.x

Replace `5.x.x` with the CDH version number. If you're using the Linux package, you can distinguish by the location of the JAR files: `/usr/lib/hadoop/client-0.20` for MRv1 APIs and `/usr/lib/hadoop/client` for YARN.



Important: Make sure that one and only one version of the `hadoop-client` artifact is available to your project. Mixing MRv1 and YARN `hadoop-client` artifacts in the same application could lead to failures that are hard to debug.

Using `hadoop-client` for Maven-based Java Projects

Make sure you add the following dependency specification to your `pom.xml` file:

```
<dependency>
  <groupId>org.apache.hadoop</groupId>
  <artifactId>hadoop-client</artifactId>
  <version>VERSION</version>
  <scope>provided</scope>
</dependency>
```

Version and Download Information

See [Table 42: CDH Version Strings for Maven](#) on page 539 for the `VERSION` string.

Using hadoop-client for Ivy-based Java Projects

Make sure you add the following dependency specification to your `ivy.xml` file:

```
<dependency org="org.apache.hadoop" name="hadoop-client" rev="VERSION"
conf="default->provided"/>
```

where the `<VERSION>` string can be either `2.2.0-cdh5.x.x` for YARN APIs or `2.2.0-mr1-cdh5.x.x` for MRv1 APIs, substituting `x` for the version number.

Using JAR Files Provided in the hadoop-client Package

Make sure you add to your project all of the JAR files provided under `/usr/lib/hadoop/client-0.20` (for MRv1 APIs) or `/usr/lib/hadoop/client` (for YARN).

For example, you can add this location to the JVM classpath:

```
$ export CLASSPATH=/usr/lib/hadoop/client-0.20/*
```

Maven Artifacts for CDH 5.16.x Releases



Important: If your cluster hosts are running CDH 5.16.x and the Ubuntu 16.04 or Debian 9 operating system and you are using the OpenJDK SDK to compile applications, note the following:

The following OpenJDK versions will cause the Maven Surefire plug-in to fail:

- Ubuntu 16.04: 8u181-b13-1ubuntu0.16.04.1
- Debian 9: 8u181-b13-2~deb9u1

with the following error message:

```
Error: Could not find or load main class
org.apache.maven.surefire.booter.ForkedBooter
```

Cloudera recommends that you use an earlier version of OpenJDK to compile your applications.

For more information, see the following bug reports:

- Debian: <https://bugs.debian.org/cgi-bin/bugreport.cgi?bug=911925>
- Ubuntu: <https://bugs.launchpad.net/ubuntu/+source/openjdk-8/+bug/1800792>

CDH 5.16.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.16.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.16.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.16.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.16.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.16.2

Project	groupId	artifactId	version
	org.apache.avro	avro-protobuf	1.7.6-cdh5.16.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.16.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.16.2
	org.apache.avro	avro-tools	1.7.6-cdh5.16.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.16.2
	org.apache.avro	trevni-core	1.7.6-cdh5.16.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.16.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.16.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.16.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.16.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.16.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.16.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.16.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.16.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.16.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.16.2
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-auth	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-config-filter-api	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-core	1.6.0-cdh5.16.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-environment-variable-config-filter	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-external-process-config-filter	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-hadoop-credential-store-config-filter	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-node	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.16.2
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.16.2
	org.apache.flume	flume-tools	1.6.0-cdh5.16.2
	org.apache.flume. flume-ng-channels	flume-file-channel	1.6.0-cdh5.16.2
	org.apache.flume. flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.16.2
	org.apache.flume. flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.16.2
	org.apache.flume. flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.16.2
	org.apache.flume. flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.16.2
	org.apache.flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.16.2
	org.apache.flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.16.2
	org.apache.flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.16.2
	org.apache.flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.16.2

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.16. 2
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.16. 2
	org.apache. flume. flume-shared	flume-shared-kafka-test	1.6.0-cdh5.16. 2
Apache Hadoop	org.apache. hadoop	hadoop-annotations	2.6.0-cdh5.16. 2
	org.apache. hadoop	hadoop-ant	2.6.0-cdh5.16. 2
	org.apache. hadoop	hadoop-archive-logs	2.6.0-cdh5.16. 2
	org.apache. hadoop	hadoop-archives	2.6.0-cdh5.16. 2
	org.apache. hadoop	hadoop-assemblies	2.6.0-cdh5.16. 2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.16.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.16.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.16.2
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.16.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.16.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.16.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.16.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.16.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.16.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.16.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.16.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-it	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.16.2

Project	groupId	artifactId	version
	org.apache.hbase	hbase-rest	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-server	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.16.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.16.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.16.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.16.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.16.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.16.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.16.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.16.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.16.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.16.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.16.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.16.2
	com.ngdata	hbase-sep-api	1.5-cdh5.16.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.16.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.16.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.16.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.16.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.16.2
	org.apache.hive	hive-ant	1.1.0-cdh5.16.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.16.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-classification	1.1.0-cdh5.16.2
	org.apache.hive	hive-cli	1.1.0-cdh5.16.2
	org.apache.hive	hive-common	1.1.0-cdh5.16.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.16.2
	org.apache.hive	hive-exec	1.1.0-cdh5.16.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.16.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.16.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.16.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.16.2
	org.apache.hive	hive-serde	1.1.0-cdh5.16.2
	org.apache.hive	hive-service	1.1.0-cdh5.16.2
	org.apache.hive	hive-shims	1.1.0-cdh5.16.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.16.2
	org.apache.hive	spark-client	1.1.0-cdh5.16.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.16.2
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.16.2
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.16.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.16.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.16.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.16.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.16.2

Project	groupId	artifactId	version
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.16.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.16.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.16.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.16.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.16.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.16.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.16.2
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.16.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.16.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.16.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.16.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.16.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.16.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.16.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.16.2
	org.kitesdk	kite-tools	1.0.0-cdh5.16.2
Apache Kudu	org.apache.kudu	kudu-client	1.7.0-cdh5.16.2
	org.apache.kudu	kudu-client-tools	1.7.0-cdh5.16.2
	org.apache.kudu	kudu-flume-sink	1.7.0-cdh5.16.2
	org.apache.kudu	kudu-mapreduce	1.7.0-cdh5.16.2
	org.apache.kudu	kudu-spark2-tools_2.11	1.7.0-cdh5.16.2
	org.apache.kudu	kudu-spark2_2.11	1.7.0-cdh5.16.2
	org.apache.kudu	kudu-spark_2.10	1.7.0-cdh5.16.2
Llama	com.cloudera.llama	llama	1.0.0-cdh5.16.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.16.2
	org.apache.mahout	mahout-core	0.9-cdh5.16.2

Project	groupId	artifactId	version
	org.apache.mahout	mahout-examples	0.9-cdh5.16.2
	org.apache.mahout	mahout-integration	0.9-cdh5.16.2
	org.apache.mahout	mahout-math	0.9-cdh5.16.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.16.2
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-core	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-examples	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.16.2.oozie-4.1.0-cdh5.16.2
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.16.2.oozie-4.1.0-cdh5.16.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.16.2.oozie-4.1.0-cdh5.16.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.16.2.oozie-4.1.0-cdh5.16.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.16.2.oozie-4.1.0-cdh5.16.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.16.2.oozie-4.1.0-cdh5.16.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.16.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.16.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.16.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.16.2
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.16.2
	com.twitter	parquet-cascading	1.5.0-cdh5.16.2
	com.twitter	parquet-column	1.5.0-cdh5.16.2
	com.twitter	parquet-common	1.5.0-cdh5.16.2
	com.twitter	parquet-encoding	1.5.0-cdh5.16.2
	com.twitter	parquet-format	2.1.0-cdh5.16.2
	com.twitter	parquet-generator	1.5.0-cdh5.16.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.16.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.16.2
	com.twitter	parquet-jackson	1.5.0-cdh5.16.2
	com.twitter	parquet-pig	1.5.0-cdh5.16.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.16.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.16.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.16.2

Project	groupId	artifactId	version
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.16.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.16.2
	com.twitter	parquet-thrift	1.5.0-cdh5.16.2
	com.twitter	parquet-tools	1.5.0-cdh5.16.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.16.2
	org.apache.pig	piggybank	0.12.0-cdh5.16.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.16.2
	org.apache.pig	pigunit	0.12.0-cdh5.16.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.16.2
	com.cloudera.search	search-mr	1.0.0-cdh5.16.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.16.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-dist	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.16.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.16.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.16.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.16.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.16.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.16.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.16.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.16.2
	org.apache.solr	solr-cell	4.10.3-cdh5.16.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.16.2
	org.apache.solr	solr-core	4.10.3-cdh5.16.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.16.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.16.2
	org.apache.solr	solr-langid	4.10.3-cdh5.16.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.16.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.16.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.16.2
	org.apache.solr	solr-uima	4.10.3-cdh5.16.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.16.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.16.2

Project	groupId	artifactId	version
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.16.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.16.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.16.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.16.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.16.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.16.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.16.2
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.16.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.16.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.16.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.16.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.16.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.16.2
	org.apache.sqoop	test	1.99.5-cdh5.16.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.16.2
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.16.2
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.16.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.16.2
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.16.2
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.16.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.16.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.16.2

Project	groupId	artifactId	version
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.16.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.16.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.16.2

Version and Download Information

Project	groupId	artifactId	version
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.16.2

CDH 5.16.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.16.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.16.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.16.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.16.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.16.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.16.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.16.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.16.1
	org.apache.avro	avro-tools	1.7.6-cdh5.16.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.16.1
	org.apache.avro	trevni-core	1.7.6-cdh5.16.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.16.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.16.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.16.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.16.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.16.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.16.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.16.1

Project	groupId	artifactId	version
	org.apache.crunch	crunch-spark	0.11.0-cdh5.16.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.16.1
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.16.1
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-config-filter-api	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-core	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-environment-variable-config-filter	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-external-process-config-filter	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-hadoop-credential-store-config-filter	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-node	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.16.1
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.16.1
	org.apache.flume	flume-tools	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.16.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sources	flume-jms-source	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.7.0-cdh5.16.1
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.7.0-cdh5.16.1

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.16.1
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.16.1
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.16.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.16.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.16.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.16.1
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.16.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.16.1
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.16.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.16.1
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.16.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.16.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.16.1
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.16.1
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-client	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-common	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-examples	1.2.0-cdh5.16.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-it	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-rest	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-server	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-shell	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-spark	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.16.1
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.16.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.16.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.16.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.16.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.16.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.16.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.16.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.16.1

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.16.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.16.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.16.1
	com.ngdata	hbase-sep-api	1.5-cdh5.16.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.16.1
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.16.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.16.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.16.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.16.1
	org.apache.hive	hive-ant	1.1.0-cdh5.16.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.16.1
	org.apache.hive	hive-classification	1.1.0-cdh5.16.1
	org.apache.hive	hive-cli	1.1.0-cdh5.16.1
	org.apache.hive	hive-common	1.1.0-cdh5.16.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.16.1
	org.apache.hive	hive-exec	1.1.0-cdh5.16.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.16.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.16.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.16.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.16.1
	org.apache.hive	hive-serde	1.1.0-cdh5.16.1
	org.apache.hive	hive-service	1.1.0-cdh5.16.1
	org.apache.hive	hive-shims	1.1.0-cdh5.16.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.16.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	spark-client	1.1.0-cdh5.16.1
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.16.1
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.16.1
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.16.1
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.16.1
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.16.1
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.16.1
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.16.1
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.16.1
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.16.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.16.1
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.16.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.16.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.16.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.16.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.16.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.16.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.16.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.16.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.16.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.16.1

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.16.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.16.1
	org.kitesdk	kite-tools	1.0.0-cdh5.16.1
Apache Kudu	org.apache.kudu	kudu-client	1.7.0-cdh5.16.1
	org.apache.kudu	kudu-client-tools	1.7.0-cdh5.16.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.kudu	kudu-flume-sink	1.7.0-cdh5.16.1
	org.apache.kudu	kudu-mapreduce	1.7.0-cdh5.16.1
	org.apache.kudu	kudu-spark2-tools_2.11	1.7.0-cdh5.16.1
	org.apache.kudu	kudu-spark2_2.11	1.7.0-cdh5.16.1
	org.apache.kudu	kudu-spark_2.10	1.7.0-cdh5.16.1
Llama	com.cloudera.llama	llama	1.0.0-cdh5.16.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.16.1
	org.apache.mahout	mahout-core	0.9-cdh5.16.1
	org.apache.mahout	mahout-examples	0.9-cdh5.16.1
	org.apache.mahout	mahout-integration	0.9-cdh5.16.1
	org.apache.mahout	mahout-math	0.9-cdh5.16.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.16.1
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.16.1
	org.apache.oozie	oozie-core	4.1.0-cdh5.16.1
	org.apache.oozie	oozie-examples	4.1.0-cdh5.16.1
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.16.1.oozie-4.1.0-cdh5.16.1
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.16.1.oozie-4.1.0-cdh5.16.1
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.16.1.oozie-4.1.0-cdh5.16.1
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.16.1

Project	groupId	artifactId	version
			1.oozie-4.1. 0-cdh5.16.1
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.16. 1.oozie-4.1. 0-cdh5.16.1
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.16. 1.oozie-4.1. 0-cdh5.16.1
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.16. 1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.16. 1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.16. 1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.16. 1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.16. 1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.16. 1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.16. 1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.16. 1
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.16. 1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.16. 1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.16. 1
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.16. 1
	com.twitter	parquet-cascading	1.5.0-cdh5.16. 1
	com.twitter	parquet-column	1.5.0-cdh5.16. 1
	com.twitter	parquet-common	1.5.0-cdh5.16. 1
	com.twitter	parquet-encoding	1.5.0-cdh5.16. 1
	com.twitter	parquet-format	2.1.0-cdh5.16. 1
	com.twitter	parquet-generator	1.5.0-cdh5.16. 1

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-hadoop	1.5.0-cdh5.16.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.16.1
	com.twitter	parquet-jackson	1.5.0-cdh5.16.1
	com.twitter	parquet-pig	1.5.0-cdh5.16.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.16.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.16.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.16.1
	com.twitter	parquet-scooage_2.10	1.5.0-cdh5.16.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.16.1
	com.twitter	parquet-thrift	1.5.0-cdh5.16.1
	com.twitter	parquet-tools	1.5.0-cdh5.16.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.16.1
	org.apache.pig	piggybank	0.12.0-cdh5.16.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.16.1
	org.apache.pig	pigunit	0.12.0-cdh5.16.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.16.1
	com.cloudera.search	search-mr	1.0.0-cdh5.16.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.16.1

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.16.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.16.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.16.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.16.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-core	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.16.1

Project	groupId	artifactId	version
	org.apache.lucene	lucene-join	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.16.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.16.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.16.1
	org.apache.solr	solr-cell	4.10.3-cdh5.16.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.16.1
	org.apache.solr	solr-core	4.10.3-cdh5.16.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.16.1
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.16.1
	org.apache.solr	solr-langid	4.10.3-cdh5.16.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.16.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.16.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.16.1
	org.apache.solr	solr-uima	4.10.3-cdh5.16.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-velocity	4.10.3-cdh5.16.1
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.16.1

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.16.1
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.16.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.16.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.16.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.16.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.16.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.16.1
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.16.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.16.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.16.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.16.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.16.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.16.1
	org.apache.sqoop	test	1.99.5-cdh5.16.1
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.16.1
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.16.1
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.16.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.16.1
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.16.1
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.16.1
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.16.1
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.16.1
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.16.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.16.1

Project	groupId	artifactId	version
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.16.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.16.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.16.1

Maven Artifacts for CDH 5.15.x Releases CDH 5.15.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.15.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.15.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.15.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.15.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.15.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.15.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.15.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.15.2
	org.apache.avro	avro-tools	1.7.6-cdh5.15.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.15.2
	org.apache.avro	trevni-core	1.7.6-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.15.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.15.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.15.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.15.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.15.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.15.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.15.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.15.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.15.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.15.2
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-config-filter-api	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-core	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-environment-variable-config-filter	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-external-process-config-filter	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-hadoop-credential-store-config-filter	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-node	1.7.0-cdh5.15.2
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.15.2

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.15.2
	org.apache.flume	flume-tools	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.15.2
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sources	flume-jms-source	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.7.0-cdh5.15.2
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.15.2
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.15.2
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.15.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.15.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.15.2
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.15.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.15.2
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.15.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.15.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.15.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.15.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.15.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.15.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-it	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-rest	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-server	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.15.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.15.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.15.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.15.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.15.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.15.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.15.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.15.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.15.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.15.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.15.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.15.2
	com.ngdata	hbase-sep-api	1.5-cdh5.15.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.15.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.15.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.15.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.15.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.15.2
	org.apache.hive	hive-ant	1.1.0-cdh5.15.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.15.2
	org.apache.hive	hive-classification	1.1.0-cdh5.15.2
	org.apache.hive	hive-cli	1.1.0-cdh5.15.2
	org.apache.hive	hive-common	1.1.0-cdh5.15.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.15.2
	org.apache.hive	hive-exec	1.1.0-cdh5.15.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.15.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.15.2

Project	groupId	artifactId	version
	org.apache.hive	hive-jdbc	1.1.0-cdh5.15.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.15.2
	org.apache.hive	hive-serde	1.1.0-cdh5.15.2
	org.apache.hive	hive-service	1.1.0-cdh5.15.2
	org.apache.hive	hive-shims	1.1.0-cdh5.15.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.15.2
	org.apache.hive	spark-client	1.1.0-cdh5.15.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.15.2
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.15.2
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.15.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.15.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.15.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.15.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.15.2
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.15.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.15.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.15.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.15.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.15.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.15.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.15.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.15.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.15.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.15.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.15.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.15.2

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.15.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.15.2
	org.kitesdk	kite-tools	1.0.0-cdh5.15.2
Apache Kudu	org.apache.kudu	kudu-client	1.7.0-cdh5.15.2
	org.apache.kudu	kudu-client-tools	1.7.0-cdh5.15.2
	org.apache.kudu	kudu-flume-sink	1.7.0-cdh5.15.2
	org.apache.kudu	kudu-mapreduce	1.7.0-cdh5.15.2
	org.apache.kudu	kudu-spark2-tools_2.11	1.7.0-cdh5.15.2
	org.apache.kudu	kudu-spark2_2.11	1.7.0-cdh5.15.2
	org.apache.kudu	kudu-spark_2.10	1.7.0-cdh5.15.2
Llama	com.cloudera.llama	llama	1.0.0-cdh5.15.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.15.2
	org.apache.mahout	mahout-core	0.9-cdh5.15.2
	org.apache.mahout	mahout-examples	0.9-cdh5.15.2
	org.apache.mahout	mahout-integration	0.9-cdh5.15.2
	org.apache.mahout	mahout-math	0.9-cdh5.15.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.15.2
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-core	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-examples	4.1.0-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.15.2.oozie-4.1.0-cdh5.15.2
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.15.2.oozie-4.1.0-cdh5.15.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.15.2.oozie-4.1.0-cdh5.15.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.15.2.oozie-4.1.0-cdh5.15.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.15.2.oozie-4.1.0-cdh5.15.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.15.2.oozie-4.1.0-cdh5.15.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.15.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.15.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.15.2

Project	groupId	artifactId	version
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.15.2
	com.twitter	parquet-cascading	1.5.0-cdh5.15.2
	com.twitter	parquet-column	1.5.0-cdh5.15.2
	com.twitter	parquet-common	1.5.0-cdh5.15.2
	com.twitter	parquet-encoding	1.5.0-cdh5.15.2
	com.twitter	parquet-format	2.1.0-cdh5.15.2
	com.twitter	parquet-generator	1.5.0-cdh5.15.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.15.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.15.2
	com.twitter	parquet-jackson	1.5.0-cdh5.15.2
	com.twitter	parquet-pig	1.5.0-cdh5.15.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.15.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.15.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.15.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.15.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.15.2
	com.twitter	parquet-thrift	1.5.0-cdh5.15.2
	com.twitter	parquet-tools	1.5.0-cdh5.15.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.15.2
	org.apache.pig	piggybank	0.12.0-cdh5.15.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.pig	pigunit	0.12.0-cdh5.15.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.15.2
	com.cloudera.search	search-mr	1.0.0-cdh5.15.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.15.2

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.15.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.15.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.15.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.15.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.15.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.15.2
	org.apache.solr	solr-cell	4.10.3-cdh5.15.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.15.2
	org.apache.solr	solr-core	4.10.3-cdh5.15.2

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.15.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.15.2
	org.apache.solr	solr-langid	4.10.3-cdh5.15.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.15.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.15.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.15.2
	org.apache.solr	solr-uima	4.10.3-cdh5.15.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.15.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.15.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.15.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.15.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.15.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.15.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.15.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.15.2
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.15.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.15.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.15.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.15.2

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.15.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.15.2
	org.apache.sqoop	test	1.99.5-cdh5.15.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.15.2
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.15.2
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.15.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.15.2
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.15.2
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.15.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.15.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.15.2
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.15.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.15.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-druid	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.15.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.15.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.15.2

CDH 5.15.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.15.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.15.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.15.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.15.1
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.15.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.15.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.15.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.15.1
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.15.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.15.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.15.1
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.15.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.15.1
	org.apache.hive	hive-ant	1.1.0-cdh5.15.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.15.1
	org.apache.hive	hive-classification	1.1.0-cdh5.15.1
	org.apache.hive	hive-cli	1.1.0-cdh5.15.1
	org.apache.hive	hive-common	1.1.0-cdh5.15.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.15.1
	org.apache.hive	hive-exec	1.1.0-cdh5.15.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.15.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.15.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.15.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.15.1
	org.apache.hive	hive-serde	1.1.0-cdh5.15.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-service	1.1.0-cdh5.15.1
	org.apache.hive	hive-shims	1.1.0-cdh5.15.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.15.1
	org.apache.hive	spark-client	1.1.0-cdh5.15.1
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.15.1
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.15.1
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.15.1
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.15.1
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.15.1
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.15.1
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.15.1
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.15.1
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.15.1
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-client	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-common	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-examples	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.15.1

Project	groupId	artifactId	version
	org.apache.hbase	hbase-it	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-rest	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-server	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-shell	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-spark	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.15.1
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.15.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.15.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.15.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.15.1
	org.apache.pig	piggybank	0.12.0-cdh5.15.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.15.1
	org.apache.pig	pigunit	0.12.0-cdh5.15.1
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.15.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-config-filter-api	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-core	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-environment-variable-config-filter	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-external-process-config-filter	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-hadoop-credential-store-config-filter	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-node	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.15.1
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.15.1
	org.apache.flume	flume-tools	1.7.0-cdh5.15.1
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.15.1
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.15.1
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.15.1
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.15.1
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.15.1
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.15.1

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.15. 1
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sources	flume-jms-source	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.7.0-cdh5.15. 1
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.7.0-cdh5.15. 1
	org.apache. flume. flume-shared	flume-shared-kafka-test	1.7.0-cdh5.15. 1

Version and Download Information

Project	groupId	artifactId	version
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-core	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-examples	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.15.1.oozie-4.1.0-cdh5.15.1
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.15.1.oozie-4.1.0-cdh5.15.1
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.15.1.oozie-4.1.0-cdh5.15.1
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.15.1.oozie-4.1.0-cdh5.15.1
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.15.1.oozie-4.1.0-cdh5.15.1
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.15.1.oozie-4.1.0-cdh5.15.1
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.15.1

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.15.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.15.1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.15.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.15.1
	org.apache.mahout	mahout-core	0.9-cdh5.15.1
	org.apache.mahout	mahout-examples	0.9-cdh5.15.1
	org.apache.mahout	mahout-integration	0.9-cdh5.15.1
	org.apache.mahout	mahout-math	0.9-cdh5.15.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.15.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.15.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.15.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.15.1
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.15.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.15.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.15.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.15.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.15.1
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.15.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.15.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.15.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.15.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.15.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.15.1
	org.apache.sqoop	test	1.99.5-cdh5.15.1
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.15.1

Project	groupId	artifactId	version
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.15.1
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.15.1
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.15.1
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.15.1
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.15.1
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.15.1
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.15.1
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.15.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.15.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.15.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.15.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.15.1
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.15.1
	com.twitter	parquet-cascading	1.5.0-cdh5.15.1

Project	groupId	artifactId	version
	com.twitter	parquet-column	1.5.0-cdh5.15.1
	com.twitter	parquet-common	1.5.0-cdh5.15.1
	com.twitter	parquet-encoding	1.5.0-cdh5.15.1
	com.twitter	parquet-format	2.1.0-cdh5.15.1
	com.twitter	parquet-generator	1.5.0-cdh5.15.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.15.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.15.1
	com.twitter	parquet-jackson	1.5.0-cdh5.15.1
	com.twitter	parquet-pig	1.5.0-cdh5.15.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.15.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.15.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.15.1
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.15.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.15.1
	com.twitter	parquet-thrift	1.5.0-cdh5.15.1
	com.twitter	parquet-tools	1.5.0-cdh5.15.1
Llama	com.cloudera.llama	llama	1.0.0-cdh5.15.1
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.15.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.15.1
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.15.1

Project	groupId	artifactId	version
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.15.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.15.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.15.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.15.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.15.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.15.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.15.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.15.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.15.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.15.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.15.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.15.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.15.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.15.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.15.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.15.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.15.1
	org.apache.avro	avro-tools	1.7.6-cdh5.15.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.15.1
	org.apache.avro	trevni-core	1.7.6-cdh5.15.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.15.1

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.15.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.15.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.15.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.15.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.15.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.15.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.15.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.15.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.15.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.15.1

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.15.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.15.1
	org.kitesdk	kite-tools	1.0.0-cdh5.15.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-core	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.15.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-join	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.15.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.15.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.15.1
	org.apache.solr	solr-cell	4.10.3-cdh5.15.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.15.1
	org.apache.solr	solr-core	4.10.3-cdh5.15.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.15.1
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.15.1
	org.apache.solr	solr-langid	4.10.3-cdh5.15.1

Project	groupId	artifactId	version
	org.apache.solr	solr-security-util	4.10.3-cdh5.15.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.15.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.15.1
	org.apache.solr	solr-uima	4.10.3-cdh5.15.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.15.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.15.1
	com.cloudera.search	search-mr	1.0.0-cdh5.15.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.15.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.15.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.15.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.15.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.15.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.15.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.15.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.15.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.15.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.15.1
	com.ngdata	hbase-sep-api	1.5-cdh5.15.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.15.1
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.15.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.15.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.15.1
Apache Kudu	org.apache.kudu	kudu-client	1.7.0-cdh5.15.1
	org.apache.kudu	kudu-client-tools	1.7.0-cdh5.15.1
	org.apache.kudu	kudu-flume-sink	1.7.0-cdh5.15.1
	org.apache.kudu	kudu-mapreduce	1.7.0-cdh5.15.1
	org.apache.kudu	kudu-spark2-tools_2.11	1.7.0-cdh5.15.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.kudu	kudu-spark2_2.11	1.7.0-cdh5.15.1
	org.apache.kudu	kudu-spark_2.10	1.7.0-cdh5.15.1

CDH 5.15.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.15.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.15.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.15.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.15.0
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.15.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.15.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.15.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.15.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.15.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.15.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.15.0
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.15.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.15.0
	org.apache.hive	hive-ant	1.1.0-cdh5.15.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.15.0
	org.apache.hive	hive-classification	1.1.0-cdh5.15.0

Project	groupId	artifactId	version
	org.apache.hive	hive-cli	1.1.0-cdh5.15.0
	org.apache.hive	hive-common	1.1.0-cdh5.15.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.15.0
	org.apache.hive	hive-exec	1.1.0-cdh5.15.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.15.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.15.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.15.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.15.0
	org.apache.hive	hive-serde	1.1.0-cdh5.15.0
	org.apache.hive	hive-service	1.1.0-cdh5.15.0
	org.apache.hive	hive-shims	1.1.0-cdh5.15.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.15.0
	org.apache.hive	spark-client	1.1.0-cdh5.15.0
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.15.0
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.15.0
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.15.0
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.15.0
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.15.0
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.15.0
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.15.0
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.15.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.15.0
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-client	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-common	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-examples	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-it	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-rest	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-server	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-shell	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-spark	1.2.0-cdh5.15.0

Project	groupId	artifactId	version
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.15.0
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.15.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.15.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.15.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.15.0
	org.apache.pig	piggybank	0.12.0-cdh5.15.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.15.0
	org.apache.pig	pigunit	0.12.0-cdh5.15.0
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-config-filter-api	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-core	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-environment-variable-config-filter	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-external-process-config-filter	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-hadoop-credential-store-config-filter	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-node	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.15.0
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.15.0
	org.apache.flume	flume-tools	1.7.0-cdh5.15.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.15.0
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.15.0
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.15.0

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sources	flume-jms-source	1.7.0-cdh5.15. 0
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.7.0-cdh5.15. 0
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.7.0-cdh5.15. 0
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.7.0-cdh5.15. 0
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.7.0-cdh5.15. 0
	org.apache. flume. flume-shared	flume-shared-kafka-test	1.7.0-cdh5.15. 0
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.15. 0
	org.apache. oozie	oozie-core	4.1.0-cdh5.15. 0
	org.apache. oozie	oozie-examples	4.1.0-cdh5.15. 0
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.15. 0.oozie-4.1. 0-cdh5.15.0
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.15. 0.oozie-4.1. 0-cdh5.15.0
	org.apache. oozie	oozie-hadoop-test	2.6. 0-mr1-cdh5.15. 0.oozie-4.1. 0-cdh5.15.0
	org.apache. oozie	oozie-hadoop-utils	2.6. 0-mr1-cdh5.15. 0.oozie-4.1. 0-cdh5.15.0
	org.apache. oozie	oozie-hbase	1.2.0-cdh5.15. 0.oozie-4.1. 0-cdh5.15.0
	org.apache. oozie	oozie-hcatalog	1.1.0-cdh5.15. 0.oozie-4.1. 0-cdh5.15.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.15.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.15.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.15.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.15.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.15.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.15.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.15.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.15.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.15.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.15.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.15.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.15.0
	org.apache.mahout	mahout-core	0.9-cdh5.15.0
	org.apache.mahout	mahout-examples	0.9-cdh5.15.0
	org.apache.mahout	mahout-integration	0.9-cdh5.15.0
	org.apache.mahout	mahout-math	0.9-cdh5.15.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.15.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.15.0

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cli	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.15.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.15.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.15.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.15.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.15.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.15.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.15.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.15.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.15.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.15.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.15.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.15.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.15.0
	org.apache.sqoop	test	1.99.5-cdh5.15.0
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.15.0
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.15.0
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.15.0
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.15.0
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.15.0
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.15.0
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.15.0
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.15.0
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.15.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.15.0

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.15.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.15.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.15.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.15.0
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.15.0
	com.twitter	parquet-cascading	1.5.0-cdh5.15.0
	com.twitter	parquet-column	1.5.0-cdh5.15.0
	com.twitter	parquet-common	1.5.0-cdh5.15.0
	com.twitter	parquet-encoding	1.5.0-cdh5.15.0
	com.twitter	parquet-format	2.1.0-cdh5.15.0
	com.twitter	parquet-generator	1.5.0-cdh5.15.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.15.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.15.0
	com.twitter	parquet-jackson	1.5.0-cdh5.15.0
	com.twitter	parquet-pig	1.5.0-cdh5.15.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.15.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.15.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.15.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.15.0

Project	groupId	artifactId	version
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.15.0
	com.twitter	parquet-thrift	1.5.0-cdh5.15.0
	com.twitter	parquet-tools	1.5.0-cdh5.15.0
Llama	com.cloudera.llama	llama	1.0.0-cdh5.15.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.15.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.15.0
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.15.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.15.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.15.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.15.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.15.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.15.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.15.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.15.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.15.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.15.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.15.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.15.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.15.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.15.0

Project	groupId	artifactId	version
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.15.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.15.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.15.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.15.0
	org.apache.avro	avro-tools	1.7.6-cdh5.15.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.15.0
	org.apache.avro	trevni-core	1.7.6-cdh5.15.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.15.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.15.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.15.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.15.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.15.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.15.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.15.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.15.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.15.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.15.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.15.0

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.15.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.15.0
	org.kitesdk	kite-tools	1.0.0-cdh5.15.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.15.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.15.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.15.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.15.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.15.0
	org.apache.solr	solr-cell	4.10.3-cdh5.15.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.15.0
	org.apache.solr	solr-core	4.10.3-cdh5.15.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.15.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.15.0
	org.apache.solr	solr-langid	4.10.3-cdh5.15.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.15.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.15.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.15.0
	org.apache.solr	solr-uima	4.10.3-cdh5.15.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.15.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.15.0
	com.cloudera.search	search-mr	1.0.0-cdh5.15.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.15.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.15.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.15.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.15.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.15.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.15.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.15.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.15.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.15.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.15.0
	com.ngdata	hbase-sep-api	1.5-cdh5.15.0

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-demo	1.5-cdh5.15.0
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.15.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.15.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.15.0
Apache Kudu	org.apache.kudu	kudu-client	1.7.0-cdh5.15.0
	org.apache.kudu	kudu-client-tools	1.7.0-cdh5.15.0
	org.apache.kudu	kudu-flume-sink	1.7.0-cdh5.15.0
	org.apache.kudu	kudu-mapreduce	1.7.0-cdh5.15.0
	org.apache.kudu	kudu-spark2-tools_2.11	1.7.0-cdh5.15.0
	org.apache.kudu	kudu-spark2_2.11	1.7.0-cdh5.15.0
	org.apache.kudu	kudu-spark_2.10	1.7.0-cdh5.15.0

Maven Artifacts for CDH 5.14.x Releases CDH 5.14.4

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.14.4

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.14.4
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.14.4
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.14.4
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.14.4
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.14.4
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.14.4
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.14.4
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.14.4
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.14.4
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.14.4
	org.apache.hive	hive-ant	1.1.0-cdh5.14.4
	org.apache.hive	hive-beeline	1.1.0-cdh5.14.4
	org.apache.hive	hive-classification	1.1.0-cdh5.14.4
	org.apache.hive	hive-cli	1.1.0-cdh5.14.4
	org.apache.hive	hive-common	1.1.0-cdh5.14.4
	org.apache.hive	hive-contrib	1.1.0-cdh5.14.4
	org.apache.hive	hive-exec	1.1.0-cdh5.14.4
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.14.4
	org.apache.hive	hive-hwi	1.1.0-cdh5.14.4
	org.apache.hive	hive-jdbc	1.1.0-cdh5.14.4
	org.apache.hive	hive-metastore	1.1.0-cdh5.14.4
	org.apache.hive	hive-serde	1.1.0-cdh5.14.4
	org.apache.hive	hive-service	1.1.0-cdh5.14.4
	org.apache.hive	hive-shims	1.1.0-cdh5.14.4
	org.apache.hive	hive-testutils	1.1.0-cdh5.14.4
	org.apache.hive	spark-client	1.1.0-cdh5.14.4

Project	groupId	artifactId	version
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.14.4
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.14.4
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.14.4
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.14.4
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.14.4
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.14.4
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.14.4
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.14.4
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.14.4
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-client	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-common	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-examples	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-it	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-rest	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-server	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-shell	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-spark	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.14.4
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.14.4
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.14.4
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.14.4
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.14.4
	org.apache.pig	piggybank	0.12.0-cdh5.14.4
	org.apache.pig	pigsmoke	0.12.0-cdh5.14.4
	org.apache.pig	pigunit	0.12.0-cdh5.14.4
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.14.4
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.14.4
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.14.4
	org.apache.flume	flume-ng-core	1.7.0-cdh5.14.4
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.14.4

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-node	1.7.0-cdh5.14.4
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.14.4
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.14.4
	org.apache.flume	flume-tools	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sources	flume-jms-source	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.7.0-cdh5.14.4
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.14.4
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.14.4
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-core	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-examples	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.14.4.oozie-4.1.0-cdh5.14.4
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.14.4.oozie-4.1.0-cdh5.14.4
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.14.4.oozie-4.1.0-cdh5.14.4

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.14.4.oozie-4.1.0-cdh5.14.4
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.14.4.oozie-4.1.0-cdh5.14.4
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.14.4.oozie-4.1.0-cdh5.14.4
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.14.4
	org.apache.oozie	oozie-tools	4.1.0-cdh5.14.4
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.14.4
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.14.4
	org.apache.mahout	mahout-core	0.9-cdh5.14.4
	org.apache.mahout	mahout-examples	0.9-cdh5.14.4
	org.apache.mahout	mahout-integration	0.9-cdh5.14.4
	org.apache.mahout	mahout-math	0.9-cdh5.14.4
	org.apache.mahout	mahout-math-scala	0.9-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-chef	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-cli	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-core	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-druid	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-examples	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-hama	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-pig	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-solr	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.14.4
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.14.4
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.14.4

Project	groupId	artifactId	version
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.14.4
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.14.4
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.14.4
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.14.4
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.14.4
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.14.4
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.14.4
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.14.4
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.14.4
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.14.4
	org.apache.sqoop	test	1.99.5-cdh5.14.4
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.14.4
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.14.4
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.14.4
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.14.4
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.14.4
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.14.4
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.14.4
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.14.4
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-dist	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.14.4

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.14.4
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.14.4
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.14.4
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.14.4
	com.twitter	parquet-cascading	1.5.0-cdh5.14.4
	com.twitter	parquet-column	1.5.0-cdh5.14.4
	com.twitter	parquet-common	1.5.0-cdh5.14.4
	com.twitter	parquet-encoding	1.5.0-cdh5.14.4
	com.twitter	parquet-format	2.1.0-cdh5.14.4
	com.twitter	parquet-generator	1.5.0-cdh5.14.4
	com.twitter	parquet-hadoop	1.5.0-cdh5.14.4
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.14.4
	com.twitter	parquet-jackson	1.5.0-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-pig	1.5.0-cdh5.14.4
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.14.4
	com.twitter	parquet-protobuf	1.5.0-cdh5.14.4
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.14.4
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.14.4
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.14.4
	com.twitter	parquet-thrift	1.5.0-cdh5.14.4
	com.twitter	parquet-tools	1.5.0-cdh5.14.4
Llama	com.cloudera.llama	llama	1.0.0-cdh5.14.4
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.14.4

Project	groupId	artifactId	version
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.14.4
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.14.4
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.14.4
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.14.4
	org.apache.crunch	crunch-core	0.11.0-cdh5.14.4
	org.apache.crunch	crunch-examples	0.11.0-cdh5.14.4
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.14.4
	org.apache.crunch	crunch-hive	0.11.0-cdh5.14.4
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.14.4
	org.apache.crunch	crunch-spark	0.11.0-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-test	0.11.0-cdh5.14.4
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.14.4
	org.apache.avro	avro-compiler	1.7.6-cdh5.14.4
	org.apache.avro	avro-ipc	1.7.6-cdh5.14.4
	org.apache.avro	avro-mapred	1.7.6-cdh5.14.4
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.14.4
	org.apache.avro	avro-protobuf	1.7.6-cdh5.14.4
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.14.4
	org.apache.avro	avro-thrift	1.7.6-cdh5.14.4
	org.apache.avro	avro-tools	1.7.6-cdh5.14.4
	org.apache.avro	trevni-avro	1.7.6-cdh5.14.4
	org.apache.avro	trevni-core	1.7.6-cdh5.14.4
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.14.4
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.14.4
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.14.4
	org.kitesdk	kite-data-hive	1.0.0-cdh5.14.4
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.14.4
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.14.4
	org.kitesdk	kite-data-s3	1.0.0-cdh5.14.4
	org.kitesdk	kite-data-spark	1.0.0-cdh5.14.4
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.14.4

Project	groupId	artifactId	version
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.14.4
	org.kitesdk	kite-minicluster	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.14.4
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.14.4
	org.kitesdk	kite-tools	1.0.0-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-classification	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-core	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-demo	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-facet	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-join	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-memory	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-misc	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-queries	4.10.3-cdh5.14.4

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.14.4
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.14.4
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.14.4
	org.apache.solr	solr-cell	4.10.3-cdh5.14.4
	org.apache.solr	solr-clustering	4.10.3-cdh5.14.4
	org.apache.solr	solr-core	4.10.3-cdh5.14.4
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.14.4
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.14.4
	org.apache.solr	solr-langid	4.10.3-cdh5.14.4
	org.apache.solr	solr-security-util	4.10.3-cdh5.14.4
	org.apache.solr	solr-solrj	4.10.3-cdh5.14.4
	org.apache.solr	solr-test-framework	4.10.3-cdh5.14.4
	org.apache.solr	solr-uima	4.10.3-cdh5.14.4
	org.apache.solr	solr-velocity	4.10.3-cdh5.14.4
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.14.4
	com.cloudera.search	search-mr	1.0.0-cdh5.14.4
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.14.4
	com.ngdata	hbase-indexer-cli	1.5-cdh5.14.4

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-common	1.5-cdh5.14.4
	com.ngdata	hbase-indexer-demo	1.5-cdh5.14.4
	com.ngdata	hbase-indexer-dist	1.5-cdh5.14.4
	com.ngdata	hbase-indexer-engine	1.5-cdh5.14.4
	com.ngdata	hbase-indexer-model	1.5-cdh5.14.4
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.14.4
	com.ngdata	hbase-indexer-mr	1.5-cdh5.14.4
	com.ngdata	hbase-indexer-server	1.5-cdh5.14.4
	com.ngdata	hbase-sep-api	1.5-cdh5.14.4
	com.ngdata	hbase-sep-demo	1.5-cdh5.14.4
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.14.4
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.14.4
	com.ngdata	hbase-sep-tools	1.5-cdh5.14.4
Apache Kudu	org.apache.kudu	kudu-client	1.6.0-cdh5.14.4
	org.apache.kudu	kudu-client-tools	1.6.0-cdh5.14.4
	org.apache.kudu	kudu-flume-sink	1.6.0-cdh5.14.4
	org.apache.kudu	kudu-mapreduce	1.6.0-cdh5.14.4
	org.apache.kudu	kudu-spark2-tools_2.11	1.6.0-cdh5.14.4
	org.apache.kudu	kudu-spark2_2.11	1.6.0-cdh5.14.4
	org.apache.kudu	kudu-spark_2.10	1.6.0-cdh5.14.4

CDH 5.14.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.14.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.14.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.14.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.14.2
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.14.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.14.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.14.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.14.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.14.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.14.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.14.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.14.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.14.2
	org.apache.hive	hive-ant	1.1.0-cdh5.14.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.14.2
	org.apache.hive	hive-classification	1.1.0-cdh5.14.2
	org.apache.hive	hive-cli	1.1.0-cdh5.14.2
	org.apache.hive	hive-common	1.1.0-cdh5.14.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.14.2
	org.apache.hive	hive-exec	1.1.0-cdh5.14.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.14.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.14.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.14.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.14.2
	org.apache.hive	hive-serde	1.1.0-cdh5.14.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-service	1.1.0-cdh5.14.2
	org.apache.hive	hive-shims	1.1.0-cdh5.14.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.14.2
	org.apache.hive	spark-client	1.1.0-cdh5.14.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.14.2
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.14.2
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.14.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.14.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.14.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.14.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.14.2
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.14.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.14.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.14.2

Project	groupId	artifactId	version
	org.apache.hbase	hbase-it	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-rest	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-server	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.14.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.14.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.14.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.14.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.14.2
	org.apache.pig	piggybank	0.12.0-cdh5.14.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.14.2
	org.apache.pig	pigunit	0.12.0-cdh5.14.2
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.14.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.14.2
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.14.2
	org.apache.flume	flume-ng-core	1.7.0-cdh5.14.2
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.14.2
	org.apache.flume	flume-ng-node	1.7.0-cdh5.14.2
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.14.2
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.14.2
	org.apache.flume	flume-tools	1.7.0-cdh5.14.2
	org.apache.flume. flume-ng-channels	flume-file-channel	1.7.0-cdh5.14.2
	org.apache.flume. flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.14.2
	org.apache.flume. flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.14.2
	org.apache.flume. flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.14.2
	org.apache.flume. flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.14.2
	org.apache.flume. flume-ng-sources	flume-avro-source	1.7.0-cdh5.14.2
	org.apache.flume. flume-ng-sources	flume-thrift-source	1.7.0-cdh5.14.2
	org.apache.flume. flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.14.2
	org.apache.flume. flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.14.2

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sources	flume-jms-source	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.7.0-cdh5.14. 2
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.7.0-cdh5.14. 2
	org.apache. flume. flume-shared	flume-shared-kafka-test	1.7.0-cdh5.14. 2
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.14. 2
	org.apache. oozie	oozie-core	4.1.0-cdh5.14. 2
	org.apache. oozie	oozie-examples	4.1.0-cdh5.14. 2
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.14. 2.oozie-4.1. 0-cdh5.14.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.14.2.oozie-4.1.0-cdh5.14.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.14.2.oozie-4.1.0-cdh5.14.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.14.2.oozie-4.1.0-cdh5.14.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.14.2.oozie-4.1.0-cdh5.14.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.14.2.oozie-4.1.0-cdh5.14.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.14.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.14.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.14.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.14.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.14.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.14.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.14.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.14.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.14.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.14.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.14.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.14.2
	org.apache.mahout	mahout-core	0.9-cdh5.14.2

Project	groupId	artifactId	version
	org.apache.mahout	mahout-examples	0.9-cdh5.14.2
	org.apache.mahout	mahout-integration	0.9-cdh5.14.2
	org.apache.mahout	mahout-math	0.9-cdh5.14.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.14.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.14.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-solr	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.14.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.14.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.14.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.14.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.14.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.14.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.14.2
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.14.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.14.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.14.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.14.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.14.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.14.2
	org.apache.sqoop	test	1.99.5-cdh5.14.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.14.2
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.14.2
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.14.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.14.2

Project	groupId	artifactId	version
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.14.2
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.14.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.14.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.14.2
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.14.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.14.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.14.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.14.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.14.2
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.14.2
	com.twitter	parquet-cascading	1.5.0-cdh5.14.2
	com.twitter	parquet-column	1.5.0-cdh5.14.2
	com.twitter	parquet-common	1.5.0-cdh5.14.2
	com.twitter	parquet-encoding	1.5.0-cdh5.14.2
	com.twitter	parquet-format	2.1.0-cdh5.14.2

Project	groupId	artifactId	version
	com.twitter	parquet-generator	1.5.0-cdh5.14.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.14.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.14.2
	com.twitter	parquet-jackson	1.5.0-cdh5.14.2
	com.twitter	parquet-pig	1.5.0-cdh5.14.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.14.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.14.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.14.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.14.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.14.2
	com.twitter	parquet-thrift	1.5.0-cdh5.14.2
	com.twitter	parquet-tools	1.5.0-cdh5.14.2
Llama	com.cloudera.llama	llama	1.0.0-cdh5.14.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.14.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.14.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.14.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.14.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.14.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.14.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.14.2

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.14.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.14.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.14.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.14.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.14.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.14.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.14.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.14.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.14.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.14.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.14.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.14.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.14.2
	org.apache.avro	avro-tools	1.7.6-cdh5.14.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.14.2
	org.apache.avro	trevni-core	1.7.6-cdh5.14.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.14.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.14.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.14.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.14.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.14.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.14.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.14.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.14.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.14.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.14.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.14.2

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.14.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.14.2
	org.kitesdk	kite-tools	1.0.0-cdh5.14.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.14.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-join	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.14.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.14.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.14.2
	org.apache.solr	solr-cell	4.10.3-cdh5.14.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.14.2
	org.apache.solr	solr-core	4.10.3-cdh5.14.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.14.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.14.2
	org.apache.solr	solr-langid	4.10.3-cdh5.14.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.14.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.14.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.14.2
	org.apache.solr	solr-uima	4.10.3-cdh5.14.2

Project	groupId	artifactId	version
	org.apache.solr	solr-velocity	4.10.3-cdh5.14.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.14.2
	com.cloudera.search	search-mr	1.0.0-cdh5.14.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.14.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.14.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.14.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.14.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.14.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.14.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.14.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.14.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.14.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.14.2
	com.ngdata	hbase-sep-api	1.5-cdh5.14.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.14.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.14.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.14.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.14.2

CDH 5.14.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.14.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.14.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.14.0
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.14.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.14.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.14.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.14.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.14.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.14.0
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.14.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.14.0
	org.apache.hive	hive-ant	1.1.0-cdh5.14.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.14.0
	org.apache.hive	hive-classification	1.1.0-cdh5.14.0
	org.apache.hive	hive-cli	1.1.0-cdh5.14.0
	org.apache.hive	hive-common	1.1.0-cdh5.14.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.14.0
	org.apache.hive	hive-exec	1.1.0-cdh5.14.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.14.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.14.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.14.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.14.0
	org.apache.hive	hive-serde	1.1.0-cdh5.14.0
	org.apache.hive	hive-service	1.1.0-cdh5.14.0
	org.apache.hive	hive-shims	1.1.0-cdh5.14.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.14.0

Project	groupId	artifactId	version
	org.apache.hive	spark-client	1.1.0-cdh5.14.0
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.14.0
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.14.0
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.14.0
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.14.0
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.14.0
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.14.0
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.14.0
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.14.0
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.14.0
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-client	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-common	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-examples	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-it	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-rest	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-server	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-shell	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-spark	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.14.0
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.14.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.14.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.14.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.14.0
	org.apache.pig	piggybank	0.12.0-cdh5.14.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.14.0
	org.apache.pig	pigunit	0.12.0-cdh5.14.0
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.14.0
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.14.0
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.14.0
	org.apache.flume	flume-ng-core	1.7.0-cdh5.14.0

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.14.0
	org.apache.flume	flume-ng-node	1.7.0-cdh5.14.0
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.14.0
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.14.0
	org.apache.flume	flume-tools	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sources	flume-jms-source	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.7.0-cdh5.14.0
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.14.0
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.14.0
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-core	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-examples	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.14.0.oozie-4.1.0-cdh5.14.0
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.14.0.oozie-4.1.0-cdh5.14.0
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.14.0.oozie-4.1.0-cdh5.14.0

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.14.0.oozie-4.1.0-cdh5.14.0
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.14.0.oozie-4.1.0-cdh5.14.0
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.14.0.oozie-4.1.0-cdh5.14.0
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.14.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.14.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.14.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.14.0
	org.apache.mahout	mahout-core	0.9-cdh5.14.0
	org.apache.mahout	mahout-examples	0.9-cdh5.14.0
	org.apache.mahout	mahout-integration	0.9-cdh5.14.0
	org.apache.mahout	mahout-math	0.9-cdh5.14.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.14.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.14.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.14.0

Project	groupId	artifactId	version
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.14.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.14.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.14.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.14.0
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.14.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.14.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.14.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.14.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.14.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.14.0
	org.apache.sqoop	test	1.99.5-cdh5.14.0
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.14.0
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.14.0
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.14.0
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.14.0
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.14.0
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.14.0
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.14.0
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.14.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.14.0

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.14.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.14.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.14.0
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.14.0
	com.twitter	parquet-cascading	1.5.0-cdh5.14.0
	com.twitter	parquet-column	1.5.0-cdh5.14.0
	com.twitter	parquet-common	1.5.0-cdh5.14.0
	com.twitter	parquet-encoding	1.5.0-cdh5.14.0
	com.twitter	parquet-format	2.1.0-cdh5.14.0
	com.twitter	parquet-generator	1.5.0-cdh5.14.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.14.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.14.0
	com.twitter	parquet-jackson	1.5.0-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-pig	1.5.0-cdh5.14.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.14.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.14.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.14.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.14.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.14.0
	com.twitter	parquet-thrift	1.5.0-cdh5.14.0
	com.twitter	parquet-tools	1.5.0-cdh5.14.0
Llama	com.cloudera.llama	llama	1.0.0-cdh5.14.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.14.0

Project	groupId	artifactId	version
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.14.0
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.14.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.14.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.14.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.14.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.14.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.14.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.14.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.14.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-test	0.11.0-cdh5.14.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.14.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.14.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.14.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.14.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.14.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.14.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.14.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.14.0
	org.apache.avro	avro-tools	1.7.6-cdh5.14.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.14.0
	org.apache.avro	trevni-core	1.7.6-cdh5.14.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.14.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.14.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.14.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.14.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.14.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.14.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.14.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.14.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.14.0

Project	groupId	artifactId	version
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.14.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.14.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.14.0
	org.kitesdk	kite-tools	1.0.0-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.14.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.14.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.14.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.14.0
	org.apache.solr	solr-cell	4.10.3-cdh5.14.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.14.0
	org.apache.solr	solr-core	4.10.3-cdh5.14.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.14.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.14.0
	org.apache.solr	solr-langid	4.10.3-cdh5.14.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.14.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.14.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.14.0
	org.apache.solr	solr-uima	4.10.3-cdh5.14.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.14.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.14.0
	com.cloudera.search	search-mr	1.0.0-cdh5.14.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.14.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.14.0

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-common	1.5-cdh5.14.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.14.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.14.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.14.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.14.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.14.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.14.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.14.0
	com.ngdata	hbase-sep-api	1.5-cdh5.14.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.14.0
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.14.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.14.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.14.0
Apache Kudu	org.apache.kudu	interface-annotations	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-client	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-client-tools	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-flume-sink	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-hive	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-mapreduce	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-parent	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-spark-tools	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-spark2-tools_2.11	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-spark2_2.11	1.6.0-cdh5.14.0
	org.apache.kudu	kudu-spark_2.10	1.6.0-cdh5.14.0

Maven Artifacts for CDH 5.13.x Releases CDH 5.13.3

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.13.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.13.3

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.13.3
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.13.3
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.13.3
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.13.3
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.13.3
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.13.3
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.13.3
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.13.3
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.13.3
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.13.3
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.13.3
	org.apache.hive	hive-ant	1.1.0-cdh5.13.3
	org.apache.hive	hive-beeline	1.1.0-cdh5.13.3
	org.apache.hive	hive-cli	1.1.0-cdh5.13.3
	org.apache.hive	hive-common	1.1.0-cdh5.13.3
	org.apache.hive	hive-contrib	1.1.0-cdh5.13.3
	org.apache.hive	hive-exec	1.1.0-cdh5.13.3
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.13.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-hwi	1.1.0-cdh5.13.3
	org.apache.hive	hive-jdbc	1.1.0-cdh5.13.3
	org.apache.hive	hive-metastore	1.1.0-cdh5.13.3
	org.apache.hive	hive-serde	1.1.0-cdh5.13.3
	org.apache.hive	hive-service	1.1.0-cdh5.13.3
	org.apache.hive	hive-shims	1.1.0-cdh5.13.3
	org.apache.hive	hive-testutils	1.1.0-cdh5.13.3
	org.apache.hive	spark-client	1.1.0-cdh5.13.3
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.13.3
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.13.3
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.13.3
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.13.3
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.13.3
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.13.3
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.13.3
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.13.3
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.13.3
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-client	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-common	1.2.0-cdh5.13.3

Project	groupId	artifactId	version
	org.apache.hbase	hbase-examples	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-it	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-rest	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-server	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-shell	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-spark	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.13.3
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.13.3
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.13.3
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.13.3
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.13.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.pig	piggybank	0.12.0-cdh5.13.3
	org.apache.pig	pigsmoke	0.12.0-cdh5.13.3
	org.apache.pig	pigunit	0.12.0-cdh5.13.3
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.13.3
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.13.3
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.13.3
	org.apache.flume	flume-ng-core	1.7.0-cdh5.13.3
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.13.3
	org.apache.flume	flume-ng-node	1.7.0-cdh5.13.3
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.13.3
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.13.3
	org.apache.flume	flume-tools	1.7.0-cdh5.13.3
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.13.3
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.13.3
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.13.3
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.13.3
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.13.3
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.13.3

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sources	flume-jms-source	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.7.0-cdh5.13. 3
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.7.0-cdh5.13. 3
	org.apache. flume. flume-shared	flume-shared-kafka-test	1.7.0-cdh5.13. 3

Version and Download Information

Project	groupId	artifactId	version
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-core	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-examples	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.13.3.oozie-4.1.0-cdh5.13.3
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.13.3.oozie-4.1.0-cdh5.13.3
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.13.3.oozie-4.1.0-cdh5.13.3
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.13.3.oozie-4.1.0-cdh5.13.3
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.13.3.oozie-4.1.0-cdh5.13.3
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.13.3.oozie-4.1.0-cdh5.13.3
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.13.3

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.13.3
	org.apache.oozie	oozie-tools	4.1.0-cdh5.13.3
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.13.3
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.13.3
	org.apache.mahout	mahout-core	0.9-cdh5.13.3
	org.apache.mahout	mahout-examples	0.9-cdh5.13.3
	org.apache.mahout	mahout-integration	0.9-cdh5.13.3
	org.apache.mahout	mahout-math	0.9-cdh5.13.3
	org.apache.mahout	mahout-math-scala	0.9-cdh5.13.3
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-chef	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-cli	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-core	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-druid	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-examples	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-hama	0.9.0-cdh5.13.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-pig	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-solr	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.13.3
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.13.3
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.13.3
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.13.3
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.13.3
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.13.3
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.13.3
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.13.3
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.13.3
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.13.3
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.13.3
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.13.3
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.13.3
	org.apache.sqoop	test	1.99.5-cdh5.13.3
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.13.3

Project	groupId	artifactId	version
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.13.3
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.13.3
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.13.3
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.13.3
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.13.3
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.13.3
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.13.3
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.13.3
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.13.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-dist	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.13.3
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.13.3
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.13.3
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.13.3
	com.twitter	parquet-cascading	1.5.0-cdh5.13.3

Project	groupId	artifactId	version
	com.twitter	parquet-column	1.5.0-cdh5.13.3
	com.twitter	parquet-common	1.5.0-cdh5.13.3
	com.twitter	parquet-encoding	1.5.0-cdh5.13.3
	com.twitter	parquet-format	2.1.0-cdh5.13.3
	com.twitter	parquet-generator	1.5.0-cdh5.13.3
	com.twitter	parquet-hadoop	1.5.0-cdh5.13.3
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.13.3
	com.twitter	parquet-jackson	1.5.0-cdh5.13.3
	com.twitter	parquet-pig	1.5.0-cdh5.13.3
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.13.3
	com.twitter	parquet-protobuf	1.5.0-cdh5.13.3
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.13.3
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.13.3
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.13.3
	com.twitter	parquet-thrift	1.5.0-cdh5.13.3
	com.twitter	parquet-tools	1.5.0-cdh5.13.3
Llama	com.cloudera.llama	llama	1.0.0-cdh5.13.3
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.13.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.13.3
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.13.3

Project	groupId	artifactId	version
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.13.3
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.13.3
	org.apache.crunch	crunch-core	0.11.0-cdh5.13.3
	org.apache.crunch	crunch-examples	0.11.0-cdh5.13.3
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.13.3
	org.apache.crunch	crunch-hive	0.11.0-cdh5.13.3
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.13.3
	org.apache.crunch	crunch-spark	0.11.0-cdh5.13.3
	org.apache.crunch	crunch-test	0.11.0-cdh5.13.3
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.13.3
	org.apache.avro	avro-compiler	1.7.6-cdh5.13.3
	org.apache.avro	avro-ipc	1.7.6-cdh5.13.3
	org.apache.avro	avro-mapred	1.7.6-cdh5.13.3
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.13.3
	org.apache.avro	avro-protobuf	1.7.6-cdh5.13.3
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.13.3
	org.apache.avro	avro-thrift	1.7.6-cdh5.13.3
	org.apache.avro	avro-tools	1.7.6-cdh5.13.3
	org.apache.avro	trevni-avro	1.7.6-cdh5.13.3
	org.apache.avro	trevni-core	1.7.6-cdh5.13.3
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.13.3

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.13.3
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.13.3
	org.kitesdk	kite-data-hive	1.0.0-cdh5.13.3
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.13.3
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.13.3
	org.kitesdk	kite-data-s3	1.0.0-cdh5.13.3
	org.kitesdk	kite-data-spark	1.0.0-cdh5.13.3
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.13.3
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.13.3
	org.kitesdk	kite-minicluster	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.13.3

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.13.3
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.13.3
	org.kitesdk	kite-tools	1.0.0-cdh5.13.3
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-classification	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-core	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-demo	4.10.3-cdh5.13.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-facet	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-join	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-memory	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-misc	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-queries	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.13.3
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.13.3
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.13.3
	org.apache.solr	solr-cell	4.10.3-cdh5.13.3
	org.apache.solr	solr-clustering	4.10.3-cdh5.13.3
	org.apache.solr	solr-core	4.10.3-cdh5.13.3
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.13.3
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.13.3
	org.apache.solr	solr-langid	4.10.3-cdh5.13.3

Project	groupId	artifactId	version
	org.apache.solr	solr-security-util	4.10.3-cdh5.13.3
	org.apache.solr	solr-solrj	4.10.3-cdh5.13.3
	org.apache.solr	solr-test-framework	4.10.3-cdh5.13.3
	org.apache.solr	solr-uima	4.10.3-cdh5.13.3
	org.apache.solr	solr-velocity	4.10.3-cdh5.13.3
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.13.3
	com.cloudera.search	search-mr	1.0.0-cdh5.13.3
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.13.3
	com.ngdata	hbase-indexer-cli	1.5-cdh5.13.3
	com.ngdata	hbase-indexer-common	1.5-cdh5.13.3
	com.ngdata	hbase-indexer-demo	1.5-cdh5.13.3
	com.ngdata	hbase-indexer-dist	1.5-cdh5.13.3
	com.ngdata	hbase-indexer-engine	1.5-cdh5.13.3
	com.ngdata	hbase-indexer-model	1.5-cdh5.13.3
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.13.3
	com.ngdata	hbase-indexer-mr	1.5-cdh5.13.3
	com.ngdata	hbase-indexer-server	1.5-cdh5.13.3
	com.ngdata	hbase-sep-api	1.5-cdh5.13.3
	com.ngdata	hbase-sep-demo	1.5-cdh5.13.3
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.13.3
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.13.3
	com.ngdata	hbase-sep-tools	1.5-cdh5.13.3

CDH 5.13.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.13.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.13.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.13.2
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.13.2
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.13.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.13.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.13.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.13.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.13.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.13.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.13.2
	org.apache.hive	hive-ant	1.1.0-cdh5.13.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.13.2
	org.apache.hive	hive-cli	1.1.0-cdh5.13.2
	org.apache.hive	hive-common	1.1.0-cdh5.13.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.13.2
	org.apache.hive	hive-exec	1.1.0-cdh5.13.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.13.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.13.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.13.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.13.2

Project	groupId	artifactId	version
	org.apache.hive	hive-serde	1.1.0-cdh5.13.2
	org.apache.hive	hive-service	1.1.0-cdh5.13.2
	org.apache.hive	hive-shims	1.1.0-cdh5.13.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.13.2
	org.apache.hive	spark-client	1.1.0-cdh5.13.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.13.2
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.13.2
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.13.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.13.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.13.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.13.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.13.2
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.13.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.13.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-it	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-rest	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-server	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.13.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.13.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.13.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.13.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.13.2
	org.apache.pig	piggybank	0.12.0-cdh5.13.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.13.2
	org.apache.pig	pigunit	0.12.0-cdh5.13.2

Project	groupId	artifactId	version
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.13.2
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.13.2
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.13.2
	org.apache.flume	flume-ng-core	1.7.0-cdh5.13.2
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.13.2
	org.apache.flume	flume-ng-node	1.7.0-cdh5.13.2
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.13.2
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.13.2
	org.apache.flume	flume-tools	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sources	flume-jms-source	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.7.0-cdh5.13.2
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.13.2
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.13.2
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-core	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-examples	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.13.2.oozie-4.1.0-cdh5.13.2

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.13.2.oozie-4.1.0-cdh5.13.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.13.2.oozie-4.1.0-cdh5.13.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.13.2.oozie-4.1.0-cdh5.13.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.13.2.oozie-4.1.0-cdh5.13.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.13.2.oozie-4.1.0-cdh5.13.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.13.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.13.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.13.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.13.2
	org.apache.mahout	mahout-core	0.9-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.mahout	mahout-examples	0.9-cdh5.13.2
	org.apache.mahout	mahout-integration	0.9-cdh5.13.2
	org.apache.mahout	mahout-math	0.9-cdh5.13.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.13.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.13.2

Project	groupId	artifactId	version
	org.apache.whirr	whirr-solr	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.13.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.13.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.13.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.13.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.13.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.13.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.13.2
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.13.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.13.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.13.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.13.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.13.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.13.2
	org.apache.sqoop	test	1.99.5-cdh5.13.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.13.2
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.13.2
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.13.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.13.2
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.13.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.13.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.13.2
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.13.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.13.2

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.13.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.13.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.13.2
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.13.2
	com.twitter	parquet-cascading	1.5.0-cdh5.13.2
	com.twitter	parquet-column	1.5.0-cdh5.13.2
	com.twitter	parquet-common	1.5.0-cdh5.13.2
	com.twitter	parquet-encoding	1.5.0-cdh5.13.2
	com.twitter	parquet-format	2.1.0-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-generator	1.5.0-cdh5.13.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.13.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.13.2
	com.twitter	parquet-jackson	1.5.0-cdh5.13.2
	com.twitter	parquet-pig	1.5.0-cdh5.13.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.13.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.13.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.13.2
	com.twitter	parquet-scooge_2.10	1.5.0-cdh5.13.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.13.2
	com.twitter	parquet-thrift	1.5.0-cdh5.13.2
	com.twitter	parquet-tools	1.5.0-cdh5.13.2
Llama	com.cloudera.llama	llama	1.0.0-cdh5.13.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.13.2

Project	groupId	artifactId	version
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.13.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.13.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.13.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.13.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.13.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.13.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.13.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.13.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.13.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.13.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.13.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.13.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.13.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.13.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.13.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.13.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.13.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.13.2
	org.apache.avro	avro-tools	1.7.6-cdh5.13.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.13.2
	org.apache.avro	trevni-core	1.7.6-cdh5.13.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.13.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.13.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.13.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.13.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.13.2

Project	groupId	artifactId	version
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.13.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.13.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.13.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.13.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.13.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.13.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.13.2
	org.kitesdk	kite-tools	1.0.0-cdh5.13.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.13.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-join	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.13.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.13.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.13.2
	org.apache.solr	solr-cell	4.10.3-cdh5.13.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.13.2
	org.apache.solr	solr-core	4.10.3-cdh5.13.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.13.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.13.2
	org.apache.solr	solr-langid	4.10.3-cdh5.13.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.13.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.13.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.13.2
	org.apache.solr	solr-uima	4.10.3-cdh5.13.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-velocity	4.10.3-cdh5.13.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.13.2
	com.cloudera.search	search-mr	1.0.0-cdh5.13.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.13.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.13.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.13.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.13.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.13.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.13.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.13.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.13.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.13.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.13.2
	com.ngdata	hbase-sep-api	1.5-cdh5.13.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.13.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.13.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.13.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.13.2
Apache Kudu	org.apache.kudu	interface-annotations	1.5.0-cdh5.13.2
	org.apache.kudu	kudu-client	1.5.0-cdh5.13.2
	org.apache.kudu	kudu-client-tools	1.5.0-cdh5.13.2
	org.apache.kudu	kudu-flume-sink	1.5.0-cdh5.13.2
	org.apache.kudu	kudu-hive	1.5.0-cdh5.13.2
	org.apache.kudu	kudu-mapreduce	1.5.0-cdh5.13.2
	org.apache.kudu	kudu-parent	1.5.0-cdh5.13.2
	org.apache.kudu	kudu-spark-tools	1.5.0-cdh5.13.2
	org.apache.kudu	kudu-spark2-tools_2.11	1.5.0-cdh5.13.2

Project	groupId	artifactId	version
	org.apache.kudu	kudu-spark2_2.11	1.5.0-cdh5.13.2
	org.apache.kudu	kudu-spark_2.10	1.5.0-cdh5.13.2

CDH 5.13.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.13.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.13.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.13.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.13.1
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.13.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.13.1
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.13.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.13.1
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.13.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.13.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.13.1
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.13.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.13.1
	org.apache.hive	hive-ant	1.1.0-cdh5.13.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.13.1
	org.apache.hive	hive-cli	1.1.0-cdh5.13.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-common	1.1.0-cdh5.13.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.13.1
	org.apache.hive	hive-exec	1.1.0-cdh5.13.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.13.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.13.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.13.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.13.1
	org.apache.hive	hive-serde	1.1.0-cdh5.13.1
	org.apache.hive	hive-service	1.1.0-cdh5.13.1
	org.apache.hive	hive-shims	1.1.0-cdh5.13.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.13.1
	org.apache.hive	spark-client	1.1.0-cdh5.13.1
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.13.1
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.13.1
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.13.1
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.13.1
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.13.1
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.13.1
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.13.1
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.13.1
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.13.1

Project	groupId	artifactId	version
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-client	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-common	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-examples	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-it	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-rest	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-server	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-shell	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-spark	1.2.0-cdh5.13.1
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.13.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.13.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.13.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.13.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.13.1
	org.apache.pig	piggybank	0.12.0-cdh5.13.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.13.1
	org.apache.pig	pigunit	0.12.0-cdh5.13.1
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.13.1
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.13.1
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.13.1
	org.apache.flume	flume-ng-core	1.7.0-cdh5.13.1
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.13.1
	org.apache.flume	flume-ng-node	1.7.0-cdh5.13.1
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.13.1
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.13.1
	org.apache.flume	flume-tools	1.7.0-cdh5.13.1
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.13.1
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.13.1
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.13.1
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.13.1

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sources	flume-avro-source	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sources	flume-jms-source	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.7.0-cdh5.13. 1
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.7.0-cdh5.13. 1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.13.1
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.13.1
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-core	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-examples	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.13.1.oozie-4.1.0-cdh5.13.1
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.13.1.oozie-4.1.0-cdh5.13.1
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.13.1.oozie-4.1.0-cdh5.13.1
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.13.1.oozie-4.1.0-cdh5.13.1
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.13.1.oozie-4.1.0-cdh5.13.1
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.13.1.oozie-4.1.0-cdh5.13.1
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.13.1

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.13.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.13.1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.13.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.13.1
	org.apache.mahout	mahout-core	0.9-cdh5.13.1
	org.apache.mahout	mahout-examples	0.9-cdh5.13.1
	org.apache.mahout	mahout-integration	0.9-cdh5.13.1
	org.apache.mahout	mahout-math	0.9-cdh5.13.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.13.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.13.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.13.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.13.1
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.13.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.13.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.13.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.13.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.13.1
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.13.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.13.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.13.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.13.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.13.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.13.1

Project	groupId	artifactId	version
	org.apache.sqoop	test	1.99.5-cdh5.13.1
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.13.1
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.13.1
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.13.1
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.13.1
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.13.1
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.13.1
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.13.1
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.13.1
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.13.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.13.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.13.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.13.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.13.1

Project	groupId	artifactId	version
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.13.1
	com.twitter	parquet-cascading	1.5.0-cdh5.13.1
	com.twitter	parquet-column	1.5.0-cdh5.13.1
	com.twitter	parquet-common	1.5.0-cdh5.13.1
	com.twitter	parquet-encoding	1.5.0-cdh5.13.1
	com.twitter	parquet-format	2.1.0-cdh5.13.1
	com.twitter	parquet-generator	1.5.0-cdh5.13.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.13.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.13.1
	com.twitter	parquet-jackson	1.5.0-cdh5.13.1
	com.twitter	parquet-pig	1.5.0-cdh5.13.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.13.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.13.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.13.1
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.13.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.13.1
	com.twitter	parquet-thrift	1.5.0-cdh5.13.1
	com.twitter	parquet-tools	1.5.0-cdh5.13.1
Llama	com.cloudera.llama	llama	1.0.0-cdh5.13.1
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.13.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.13.1

Project	groupId	artifactId	version
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.13.1
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.13.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.13.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.13.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.13.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.13.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.13.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.13.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.13.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.13.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.13.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.13.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.13.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.13.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.13.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.13.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.13.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.13.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.13.1
	org.apache.avro	avro-tools	1.7.6-cdh5.13.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.13.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	trevni-core	1.7.6-cdh5.13.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.13.1
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.13.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.13.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.13.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.13.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.13.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.13.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.13.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.13.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.13.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.13.1

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.13.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.13.1
	org.kitesdk	kite-tools	1.0.0-cdh5.13.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.13.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-core	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-join	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.13.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.13.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.13.1
	org.apache.solr	solr-cell	4.10.3-cdh5.13.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.13.1
	org.apache.solr	solr-core	4.10.3-cdh5.13.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.13.1

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.13.1
	org.apache.solr	solr-langid	4.10.3-cdh5.13.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.13.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.13.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.13.1
	org.apache.solr	solr-uima	4.10.3-cdh5.13.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.13.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.13.1
	com.cloudera.search	search-mr	1.0.0-cdh5.13.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.13.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.13.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.13.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.13.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.13.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.13.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.13.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.13.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.13.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.13.1
	com.ngdata	hbase-sep-api	1.5-cdh5.13.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.13.1
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.13.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.13.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.13.1
Apache Kudu	org.apache.kudu	interface-annotations	1.5.0-cdh5.13.1
	org.apache.kudu	kudu-client	1.5.0-cdh5.13.1
	org.apache.kudu	kudu-client-tools	1.5.0-cdh5.13.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.kudu	kudu-flume-sink	1.5.0-cdh5.13.1
	org.apache.kudu	kudu-hive	1.5.0-cdh5.13.1
	org.apache.kudu	kudu-mapreduce	1.5.0-cdh5.13.1
	org.apache.kudu	kudu-parent	1.5.0-cdh5.13.1
	org.apache.kudu	kudu-spark-tools	1.5.0-cdh5.13.1
	org.apache.kudu	kudu-spark2-tools_2.11	1.5.0-cdh5.13.1
	org.apache.kudu	kudu-spark2_2.11	1.5.0-cdh5.13.1
	org.apache.kudu	kudu-spark_2.10	1.5.0-cdh5.13.1

CDH 5.13.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.13.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.13.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.13.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.13.0
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.13.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.13.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.13.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.13.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.13.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.13.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.13.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.13.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.13.0
	org.apache.hive	hive-ant	1.1.0-cdh5.13.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.13.0
	org.apache.hive	hive-cli	1.1.0-cdh5.13.0
	org.apache.hive	hive-common	1.1.0-cdh5.13.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.13.0
	org.apache.hive	hive-exec	1.1.0-cdh5.13.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.13.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.13.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.13.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.13.0
	org.apache.hive	hive-serde	1.1.0-cdh5.13.0
	org.apache.hive	hive-service	1.1.0-cdh5.13.0
	org.apache.hive	hive-shims	1.1.0-cdh5.13.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.13.0
	org.apache.hive	spark-client	1.1.0-cdh5.13.0
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.13.0
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.13.0
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.13.0
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.13.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.13.0
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.13.0
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.13.0
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.13.0
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.13.0
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-client	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-common	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-examples	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-it	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-rest	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-server	1.2.0-cdh5.13.0

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-shell	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-spark	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.13.0
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.13.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.13.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.13.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.13.0
	org.apache.pig	piggybank	0.12.0-cdh5.13.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.13.0
	org.apache.pig	pigunit	0.12.0-cdh5.13.0
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.13.0
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.13.0
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.13.0
	org.apache.flume	flume-ng-core	1.7.0-cdh5.13.0
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.13.0
	org.apache.flume	flume-ng-node	1.7.0-cdh5.13.0
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.13.0
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.13.0
	org.apache.flume	flume-tools	1.7.0-cdh5.13.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.13.0
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.13.0

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sources	flume-jms-source	1.7.0-cdh5.13. 0
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.7.0-cdh5.13. 0
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.7.0-cdh5.13. 0
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.7.0-cdh5.13. 0
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.7.0-cdh5.13. 0
	org.apache. flume. flume-shared	flume-shared-kafka-test	1.7.0-cdh5.13. 0
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.13. 0
	org.apache. oozie	oozie-core	4.1.0-cdh5.13. 0
	org.apache. oozie	oozie-examples	4.1.0-cdh5.13. 0
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.13. 0.oozie-4.1. 0-cdh5.13.0
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.13. 0.oozie-4.1. 0-cdh5.13.0
	org.apache. oozie	oozie-hadoop-test	2.6. 0-mr1-cdh5.13. 0.oozie-4.1. 0-cdh5.13.0
	org.apache. oozie	oozie-hadoop-utils	2.6. 0-mr1-cdh5.13. 0.oozie-4.1. 0-cdh5.13.0
	org.apache. oozie	oozie-hbase	1.2.0-cdh5.13. 0.oozie-4.1. 0-cdh5.13.0
	org.apache. oozie	oozie-hcatalog	1.1.0-cdh5.13. 0.oozie-4.1. 0-cdh5.13.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.13.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.13.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.13.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.13.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.13.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.13.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.13.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.13.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.13.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.13.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.13.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.13.0
	org.apache.mahout	mahout-core	0.9-cdh5.13.0
	org.apache.mahout	mahout-examples	0.9-cdh5.13.0
	org.apache.mahout	mahout-integration	0.9-cdh5.13.0
	org.apache.mahout	mahout-math	0.9-cdh5.13.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.13.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.13.0

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cli	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.13.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.13.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.13.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.13.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.13.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.13.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.13.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.13.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.13.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.13.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.13.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.13.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.13.0
	org.apache.sqoop	test	1.99.5-cdh5.13.0
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.13.0
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.13.0
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.13.0
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.13.0
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.13.0
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.13.0
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.13.0
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.13.0
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.13.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.13.0

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-hive-conf	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-binding-hive-follower	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.13.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.13.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.13.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.13.0
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.13.0
	com.twitter	parquet-cascading	1.5.0-cdh5.13.0
	com.twitter	parquet-column	1.5.0-cdh5.13.0
	com.twitter	parquet-common	1.5.0-cdh5.13.0
	com.twitter	parquet-encoding	1.5.0-cdh5.13.0
	com.twitter	parquet-format	2.1.0-cdh5.13.0
	com.twitter	parquet-generator	1.5.0-cdh5.13.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.13.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.13.0
	com.twitter	parquet-jackson	1.5.0-cdh5.13.0
	com.twitter	parquet-pig	1.5.0-cdh5.13.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.13.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.13.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.13.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.13.0

Project	groupId	artifactId	version
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.13.0
	com.twitter	parquet-thrift	1.5.0-cdh5.13.0
	com.twitter	parquet-tools	1.5.0-cdh5.13.0
Llama	com.cloudera.llama	llama	1.0.0-cdh5.13.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.13.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.13.0
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.13.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.13.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.13.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.13.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.13.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.13.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.13.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.13.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.13.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.13.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.13.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.13.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.13.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.13.0

Project	groupId	artifactId	version
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.13.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.13.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.13.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.13.0
	org.apache.avro	avro-tools	1.7.6-cdh5.13.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.13.0
	org.apache.avro	trevni-core	1.7.6-cdh5.13.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.13.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.13.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.13.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.13.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.13.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.13.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.13.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.13.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.13.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.13.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.13.0

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.13.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.13.0
	org.kitesdk	kite-tools	1.0.0-cdh5.13.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.13.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.13.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.13.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.13.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.13.0
	org.apache.solr	solr-cell	4.10.3-cdh5.13.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.13.0
	org.apache.solr	solr-core	4.10.3-cdh5.13.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.13.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.13.0
	org.apache.solr	solr-langid	4.10.3-cdh5.13.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.13.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.13.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.13.0
	org.apache.solr	solr-uima	4.10.3-cdh5.13.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.13.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.13.0
	com.cloudera.search	search-mr	1.0.0-cdh5.13.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.13.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.13.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.13.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.13.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.13.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.13.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.13.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.13.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.13.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.13.0
	com.ngdata	hbase-sep-api	1.5-cdh5.13.0

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-demo	1.5-cdh5.13.0
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.13.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.13.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.13.0
Apache Kudu	org.apache.kudu	interface-annotations	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-client	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-client-tools	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-flume-sink	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-hive	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-mapreduce	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-parent	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-spark-tools	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-spark2-tools_2.11	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-spark2_2.11	1.5.0-cdh5.13.0
	org.apache.kudu	kudu-spark_2.10	1.5.0-cdh5.13.0

Maven Artifacts for CDH 5.12.x Releases CDH 5.12.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.12.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.12.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.12.2
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.12.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.12.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.12.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.12.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.12.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.12.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.12.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.12.2
	org.apache.hive	hive-ant	1.1.0-cdh5.12.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.12.2
	org.apache.hive	hive-cli	1.1.0-cdh5.12.2
	org.apache.hive	hive-common	1.1.0-cdh5.12.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.12.2
	org.apache.hive	hive-exec	1.1.0-cdh5.12.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.12.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.12.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.12.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.12.2
	org.apache.hive	hive-serde	1.1.0-cdh5.12.2
	org.apache.hive	hive-service	1.1.0-cdh5.12.2

Project	groupId	artifactId	version
	org.apache.hive	hive-shims	1.1.0-cdh5.12.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.12.2
	org.apache.hive	spark-client	1.1.0-cdh5.12.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.12.2
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.12.2
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.12.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.12.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.12.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.12.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.12.2
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.12.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.12.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-it	1.2.0-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-rest	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-server	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.12.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.12.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.12.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.12.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.12.2
	org.apache.pig	piggybank	0.12.0-cdh5.12.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.12.2
	org.apache.pig	pigunit	0.12.0-cdh5.12.2
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.12.2
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.12.2

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.12.2
	org.apache.flume	flume-ng-core	1.7.0-cdh5.12.2
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.12.2
	org.apache.flume	flume-ng-node	1.7.0-cdh5.12.2
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.12.2
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.12.2
	org.apache.flume	flume-tools	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sources	flume-jms-source	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.7.0-cdh5.12.2
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.12.2
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.12.2
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-core	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-examples	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.12.2.oozie-4.1.0-cdh5.12.2
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.12.2

Project	groupId	artifactId	version
			2.oozie-4.1.0-cdh5.12.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.12.2.oozie-4.1.0-cdh5.12.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.12.2.oozie-4.1.0-cdh5.12.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.12.2.oozie-4.1.0-cdh5.12.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.12.2.oozie-4.1.0-cdh5.12.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.12.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.12.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.12.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.12.2
	org.apache.mahout	mahout-core	0.9-cdh5.12.2
	org.apache.mahout	mahout-examples	0.9-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.mahout	mahout-integration	0.9-cdh5.12.2
	org.apache.mahout	mahout-math	0.9-cdh5.12.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.12.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.12.2

Project	groupId	artifactId	version
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.12.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.12.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.12.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.12.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.12.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.12.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.12.2
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.12.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.12.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.12.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.12.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.12.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.12.2
	org.apache.sqoop	test	1.99.5-cdh5.12.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.12.2
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.12.2
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.12.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.12.2
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.12.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.12.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.12.2
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.12.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.12.2

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.12.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.12.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.12.2
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.12.2
	com.twitter	parquet-cascading	1.5.0-cdh5.12.2
	com.twitter	parquet-column	1.5.0-cdh5.12.2
	com.twitter	parquet-common	1.5.0-cdh5.12.2
	com.twitter	parquet-encoding	1.5.0-cdh5.12.2
	com.twitter	parquet-format	2.1.0-cdh5.12.2
	com.twitter	parquet-generator	1.5.0-cdh5.12.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.12.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.12.2
	com.twitter	parquet-jackson	1.5.0-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-pig	1.5.0-cdh5.12.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.12.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.12.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.12.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.12.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.12.2
	com.twitter	parquet-thrift	1.5.0-cdh5.12.2
	com.twitter	parquet-tools	1.5.0-cdh5.12.2
Llama	com.cloudera.llama	llama	1.0.0-cdh5.12.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.12.2

Project	groupId	artifactId	version
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.12.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.12.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.12.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.12.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.12.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.12.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.12.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.12.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.12.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-test	0.11.0-cdh5.12.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.12.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.12.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.12.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.12.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.12.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.12.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.12.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.12.2
	org.apache.avro	avro-tools	1.7.6-cdh5.12.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.12.2
	org.apache.avro	trevni-core	1.7.6-cdh5.12.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.12.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.12.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.12.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.12.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.12.2
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.12.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.12.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.12.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.12.2

Project	groupId	artifactId	version
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.12.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.12.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.12.2
	org.kitesdk	kite-tools	1.0.0-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.12.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.12.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.12.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.12.2
	org.apache.solr	solr-cell	4.10.3-cdh5.12.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.12.2
	org.apache.solr	solr-core	4.10.3-cdh5.12.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.12.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.12.2
	org.apache.solr	solr-langid	4.10.3-cdh5.12.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.12.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.12.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.12.2
	org.apache.solr	solr-uima	4.10.3-cdh5.12.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.12.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.12.2
	com.cloudera.search	search-mr	1.0.0-cdh5.12.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.12.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.12.2

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-common	1.5-cdh5.12.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.12.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.12.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.12.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.12.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.12.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.12.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.12.2
	com.ngdata	hbase-sep-api	1.5-cdh5.12.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.12.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.12.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.12.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.12.2

CDH 5.12.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.12.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.12.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.12.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.12.1
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.12.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.12.1
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.12.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.12.1
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.12.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.12.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.12.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.12.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.12.1
	org.apache.hive	hive-ant	1.1.0-cdh5.12.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.12.1
	org.apache.hive	hive-cli	1.1.0-cdh5.12.1
	org.apache.hive	hive-common	1.1.0-cdh5.12.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.12.1
	org.apache.hive	hive-exec	1.1.0-cdh5.12.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.12.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.12.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.12.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.12.1
	org.apache.hive	hive-serde	1.1.0-cdh5.12.1
	org.apache.hive	hive-service	1.1.0-cdh5.12.1
	org.apache.hive	hive-shims	1.1.0-cdh5.12.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.12.1
	org.apache.hive	spark-client	1.1.0-cdh5.12.1
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.12.1
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.12.1
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.12.1
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.12.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.12.1
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.12.1
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.12.1
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.12.1
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.12.1
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-client	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-common	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-examples	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-it	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-rest	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-server	1.2.0-cdh5.12.1

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-shell	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-spark	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.12.1
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.12.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.12.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.12.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.12.1
	org.apache.pig	piggybank	0.12.0-cdh5.12.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.12.1
	org.apache.pig	pigunit	0.12.0-cdh5.12.1
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.12.1
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.12.1
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.12.1
	org.apache.flume	flume-ng-core	1.7.0-cdh5.12.1
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.12.1
	org.apache.flume	flume-ng-node	1.7.0-cdh5.12.1
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.12.1
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.12.1
	org.apache.flume	flume-tools	1.7.0-cdh5.12.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.12.1
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.12.1

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sources	flume-jms-source	1.7.0-cdh5.12. 1
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.7.0-cdh5.12. 1
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.7.0-cdh5.12. 1
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.7.0-cdh5.12. 1
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.7.0-cdh5.12. 1
	org.apache. flume. flume-shared	flume-shared-kafka-test	1.7.0-cdh5.12. 1
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.12. 1
	org.apache. oozie	oozie-core	4.1.0-cdh5.12. 1
	org.apache. oozie	oozie-examples	4.1.0-cdh5.12. 1
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.12. 1.oozie-4.1. 0-cdh5.12.1
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.12. 1.oozie-4.1. 0-cdh5.12.1
	org.apache. oozie	oozie-hadoop-test	2.6. 0-mr1-cdh5.12. 1.oozie-4.1. 0-cdh5.12.1
	org.apache. oozie	oozie-hadoop-utils	2.6. 0-mr1-cdh5.12. 1.oozie-4.1. 0-cdh5.12.1
	org.apache. oozie	oozie-hbase	1.2.0-cdh5.12. 1.oozie-4.1. 0-cdh5.12.1
	org.apache. oozie	oozie-hcatalog	1.1.0-cdh5.12. 1.oozie-4.1. 0-cdh5.12.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.12.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.12.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.12.1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.12.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.12.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.12.1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.12.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.12.1
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.12.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.12.1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.12.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.12.1
	org.apache.mahout	mahout-core	0.9-cdh5.12.1
	org.apache.mahout	mahout-examples	0.9-cdh5.12.1
	org.apache.mahout	mahout-integration	0.9-cdh5.12.1
	org.apache.mahout	mahout-math	0.9-cdh5.12.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.12.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.12.1

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cli	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.12.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.12.1
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.12.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.12.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.12.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.12.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.12.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.12.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.12.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.12.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.12.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.12.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.12.1
	org.apache.sqoop	test	1.99.5-cdh5.12.1
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.12.1
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.12.1
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.12.1
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.12.1
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.12.1
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.12.1
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.12.1
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.12.1
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.12.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.12.1

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.12.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.12.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.12.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.12.1
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.12.1
	com.twitter	parquet-cascading	1.5.0-cdh5.12.1
	com.twitter	parquet-column	1.5.0-cdh5.12.1
	com.twitter	parquet-common	1.5.0-cdh5.12.1
	com.twitter	parquet-encoding	1.5.0-cdh5.12.1
	com.twitter	parquet-format	2.1.0-cdh5.12.1
	com.twitter	parquet-generator	1.5.0-cdh5.12.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.12.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.12.1
	com.twitter	parquet-jackson	1.5.0-cdh5.12.1
	com.twitter	parquet-pig	1.5.0-cdh5.12.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.12.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.12.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.12.1
	com.twitter	parquet-scooage_2.10	1.5.0-cdh5.12.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.12.1
	com.twitter	parquet-thrift	1.5.0-cdh5.12.1

Project	groupId	artifactId	version
	com.twitter	parquet-tools	1.5.0-cdh5.12.1
Llama	com.cloudera.llama	llama	1.0.0-cdh5.12.1
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.12.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.12.1
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.12.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.12.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.12.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.12.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.12.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.12.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.12.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.12.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.12.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.12.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.12.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.12.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.12.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.12.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.12.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.12.1

Project	groupId	artifactId	version
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.12.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.12.1
	org.apache.avro	avro-tools	1.7.6-cdh5.12.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.12.1
	org.apache.avro	trevni-core	1.7.6-cdh5.12.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.12.1
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.12.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.12.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.12.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.12.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.12.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.12.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.12.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.12.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.12.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.12.1

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.12.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.12.1
	org.kitesdk	kite-tools	1.0.0-cdh5.12.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.12.1

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-core	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-join	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.12.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.12.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.12.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-cell	4.10.3-cdh5.12.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.12.1
	org.apache.solr	solr-core	4.10.3-cdh5.12.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.12.1
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.12.1
	org.apache.solr	solr-langid	4.10.3-cdh5.12.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.12.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.12.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.12.1
	org.apache.solr	solr-uima	4.10.3-cdh5.12.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.12.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.12.1
	com.cloudera.search	search-mr	1.0.0-cdh5.12.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.12.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.12.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.12.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.12.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.12.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.12.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.12.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.12.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.12.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.12.1
	com.ngdata	hbase-sep-api	1.5-cdh5.12.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.12.1
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.12.1

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.12.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.12.1

CDH 5.12.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.12.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.12.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.12.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.12.0
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.12.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.12.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.12.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.12.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.12.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.12.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.12.0
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.12.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.12.0
	org.apache.hive	hive-ant	1.1.0-cdh5.12.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.12.0
	org.apache.hive	hive-cli	1.1.0-cdh5.12.0
	org.apache.hive	hive-common	1.1.0-cdh5.12.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-contrib	1.1.0-cdh5.12.0
	org.apache.hive	hive-exec	1.1.0-cdh5.12.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.12.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.12.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.12.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.12.0
	org.apache.hive	hive-serde	1.1.0-cdh5.12.0
	org.apache.hive	hive-service	1.1.0-cdh5.12.0
	org.apache.hive	hive-shims	1.1.0-cdh5.12.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.12.0
	org.apache.hive	spark-client	1.1.0-cdh5.12.0
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.12.0
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.12.0
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.12.0
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.12.0
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.12.0
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.12.0
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.12.0
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.12.0
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.12.0
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.12.0

Project	groupId	artifactId	version
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-client	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-common	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-examples	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-it	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-rest	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-server	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-shell	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-spark	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.12.0
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.12.0

Version and Download Information

Project	groupId	artifactId	version
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.12.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.12.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.12.0
	org.apache.pig	piggybank	0.12.0-cdh5.12.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.12.0
	org.apache.pig	pigunit	0.12.0-cdh5.12.0
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.12.0
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.12.0
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.12.0
	org.apache.flume	flume-ng-core	1.7.0-cdh5.12.0
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.12.0
	org.apache.flume	flume-ng-node	1.7.0-cdh5.12.0
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.12.0
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.12.0
	org.apache.flume	flume-tools	1.7.0-cdh5.12.0
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.12.0
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.12.0
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.12.0
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.12.0

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sources	flume-avro-source	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sources	flume-jms-source	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.7.0-cdh5.12. 0
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.7.0-cdh5.12. 0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.12.0
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.12.0
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-core	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-examples	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.12.0.oozie-4.1.0-cdh5.12.0
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.12.0.oozie-4.1.0-cdh5.12.0
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.12.0.oozie-4.1.0-cdh5.12.0
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.12.0.oozie-4.1.0-cdh5.12.0
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.12.0.oozie-4.1.0-cdh5.12.0
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.12.0.oozie-4.1.0-cdh5.12.0
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.12.0

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.12.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.12.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.12.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.12.0
	org.apache.mahout	mahout-core	0.9-cdh5.12.0
	org.apache.mahout	mahout-examples	0.9-cdh5.12.0
	org.apache.mahout	mahout-integration	0.9-cdh5.12.0
	org.apache.mahout	mahout-math	0.9-cdh5.12.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.12.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.12.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.12.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.12.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.12.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.12.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.12.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.12.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.12.0
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.12.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.12.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.12.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.12.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.12.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.12.0

Project	groupId	artifactId	version
	org.apache.sqoop	test	1.99.5-cdh5.12.0
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.12.0
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.12.0
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.12.0
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.12.0
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.12.0
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.12.0
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.12.0
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.12.0
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.12.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.12.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.12.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.12.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.12.0
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.12.0
	com.twitter	parquet-cascading	1.5.0-cdh5.12.0

Project	groupId	artifactId	version
	com.twitter	parquet-column	1.5.0-cdh5.12.0
	com.twitter	parquet-common	1.5.0-cdh5.12.0
	com.twitter	parquet-encoding	1.5.0-cdh5.12.0
	com.twitter	parquet-format	2.1.0-cdh5.12.0
	com.twitter	parquet-generator	1.5.0-cdh5.12.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.12.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.12.0
	com.twitter	parquet-jackson	1.5.0-cdh5.12.0
	com.twitter	parquet-pig	1.5.0-cdh5.12.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.12.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.12.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.12.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.12.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.12.0
	com.twitter	parquet-thrift	1.5.0-cdh5.12.0
	com.twitter	parquet-tools	1.5.0-cdh5.12.0
Llama	com.cloudera.llama	llama	1.0.0-cdh5.12.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.12.0
	com.databricks	spark-avro_2.10	1.1.0-cdh5.12.0
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.12.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.12.0
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.12.0

Project	groupId	artifactId	version
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.12.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.12.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.12.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.12.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.12.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.12.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.12.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.12.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.12.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.12.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.12.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.12.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.12.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.12.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.12.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.12.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.12.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.12.0
	org.apache.avro	avro-tools	1.7.6-cdh5.12.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.12.0
	org.apache.avro	trevni-core	1.7.6-cdh5.12.0

Version and Download Information

Project	groupId	artifactId	version
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.12.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.12.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.12.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.12.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.12.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.12.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.12.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.12.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.12.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.12.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.12.0

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.12.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.12.0
	org.kitesdk	kite-tools	1.0.0-cdh5.12.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.12.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-demo	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.12.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.12.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.12.0
	org.apache.solr	solr-cell	4.10.3-cdh5.12.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.12.0
	org.apache.solr	solr-core	4.10.3-cdh5.12.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.12.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.12.0

Project	groupId	artifactId	version
	org.apache.solr	solr-langid	4.10.3-cdh5.12.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.12.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.12.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.12.0
	org.apache.solr	solr-uima	4.10.3-cdh5.12.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.12.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.12.0
	com.cloudera.search	search-mr	1.0.0-cdh5.12.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.12.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.12.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.12.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.12.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.12.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.12.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.12.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.12.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.12.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.12.0
	com.ngdata	hbase-sep-api	1.5-cdh5.12.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.12.0
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.12.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.12.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.12.0

Maven Artifacts for CDH 5.11.x Releases

CDH 5.11.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.11.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-mapreduce-client-native-task	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.11.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.11.2
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.11.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.11.2
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.11.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.11.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.11.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.11.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.11.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.11.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.11.2
	org.apache.hive	hive-ant	1.1.0-cdh5.11.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.11.2
	org.apache.hive	hive-cli	1.1.0-cdh5.11.2
	org.apache.hive	hive-common	1.1.0-cdh5.11.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.11.2
	org.apache.hive	hive-exec	1.1.0-cdh5.11.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.11.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.11.2

Project	groupId	artifactId	version
	org.apache.hive	hive-jdbc	1.1.0-cdh5.11.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.11.2
	org.apache.hive	hive-serde	1.1.0-cdh5.11.2
	org.apache.hive	hive-service	1.1.0-cdh5.11.2
	org.apache.hive	hive-shims	1.1.0-cdh5.11.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.11.2
	org.apache.hive	spark-client	1.1.0-cdh5.11.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.11.2
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.11.2
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.11.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.11.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.11.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.11.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.11.2
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.11.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.11.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-it	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-rest	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-server	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.11.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.11.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.11.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.11.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.11.2
	org.apache.pig	piggybank	0.12.0-cdh5.11.2

Project	groupId	artifactId	version
	org.apache.pig	pigsmoke	0.12.0-cdh5.11.2
	org.apache.pig	pigunit	0.12.0-cdh5.11.2
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.11.2
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.11.2
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.11.2
	org.apache.flume	flume-ng-core	1.7.0-cdh5.11.2
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.11.2
	org.apache.flume	flume-ng-node	1.7.0-cdh5.11.2
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.11.2
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.11.2
	org.apache.flume	flume-tools	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sources	flume-jms-source	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.7.0-cdh5.11.2
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.11.2
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.11.2
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-core	4.1.0-cdh5.11.2

Project	groupId	artifactId	version
	org.apache.oozie	oozie-examples	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.11.2.oozie-4.1.0-cdh5.11.2
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.11.2.oozie-4.1.0-cdh5.11.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.11.2.oozie-4.1.0-cdh5.11.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.11.2.oozie-4.1.0-cdh5.11.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.11.2.oozie-4.1.0-cdh5.11.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.11.2.oozie-4.1.0-cdh5.11.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.11.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.11.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.11.2
	org.apache.mahout	mahout-core	0.9-cdh5.11.2
	org.apache.mahout	mahout-examples	0.9-cdh5.11.2
	org.apache.mahout	mahout-integration	0.9-cdh5.11.2
	org.apache.mahout	mahout-math	0.9-cdh5.11.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.11.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.11.2

Project	groupId	artifactId	version
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.11.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.11.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.11.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.11.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.11.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.11.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.11.2
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.11.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.11.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.11.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.11.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.11.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.11.2
	org.apache.sqoop	test	1.99.5-cdh5.11.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.11.2
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.11.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.11.2
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.11.2
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.11.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.11.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.11.2
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.11.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.11.2

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.11.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.11.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.11.2
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.11.2
	com.twitter	parquet-cascading	1.5.0-cdh5.11.2
	com.twitter	parquet-column	1.5.0-cdh5.11.2
	com.twitter	parquet-common	1.5.0-cdh5.11.2
	com.twitter	parquet-encoding	1.5.0-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-format	2.1.0-cdh5.11.2
	com.twitter	parquet-generator	1.5.0-cdh5.11.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.11.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.11.2
	com.twitter	parquet-jackson	1.5.0-cdh5.11.2
	com.twitter	parquet-pig	1.5.0-cdh5.11.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.11.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.11.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.11.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.11.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.11.2
	com.twitter	parquet-thrift	1.5.0-cdh5.11.2
	com.twitter	parquet-tools	1.5.0-cdh5.11.2
Llama	com.cloudera.llama	llama	1.0.0-cdh5.11.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.11.2
	com.databricks	spark-avro_2.10	1.1.0-cdh5.11.2
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.11.2

Project	groupId	artifactId	version
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.11.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.11.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.11.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-core	0.11.0-cdh5.11.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.11.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.11.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.11.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.11.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.11.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.11.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.11.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.11.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.11.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.11.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.11.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.11.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.11.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.11.2
	org.apache.avro	avro-tools	1.7.6-cdh5.11.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.11.2
	org.apache.avro	trevni-core	1.7.6-cdh5.11.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.11.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.11.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.11.2

Project	groupId	artifactId	version
	org.kitesdk	kite-data-hive	1.0.0-cdh5.11.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.11.2
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.11.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.11.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.11.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.11.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.11.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.11.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.11.2
	org.kitesdk	kite-tools	1.0.0-cdh5.11.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.11.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.11.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.11.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.11.2
	org.apache.solr	solr-cell	4.10.3-cdh5.11.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.11.2
	org.apache.solr	solr-core	4.10.3-cdh5.11.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.11.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.11.2
	org.apache.solr	solr-langid	4.10.3-cdh5.11.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.11.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.11.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-test-framework	4.10.3-cdh5.11.2
	org.apache.solr	solr-uima	4.10.3-cdh5.11.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.11.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.11.2
	com.cloudera.search	search-mr	1.0.0-cdh5.11.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.11.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.11.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.11.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.11.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.11.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.11.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.11.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.11.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.11.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.11.2
	com.ngdata	hbase-sep-api	1.5-cdh5.11.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.11.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.11.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.11.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.11.2

CDH 5.11.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.11.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.11.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.11.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.11.1
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.11.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.11.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.11.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.11.1
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.11.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.11.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.11.1
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.11.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.11.1
	org.apache.hive	hive-ant	1.1.0-cdh5.11.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.11.1
	org.apache.hive	hive-cli	1.1.0-cdh5.11.1
	org.apache.hive	hive-common	1.1.0-cdh5.11.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.11.1
	org.apache.hive	hive-exec	1.1.0-cdh5.11.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.11.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.11.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.11.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.11.1
	org.apache.hive	hive-serde	1.1.0-cdh5.11.1
	org.apache.hive	hive-service	1.1.0-cdh5.11.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-shims	1.1.0-cdh5.11.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.11.1
	org.apache.hive	spark-client	1.1.0-cdh5.11.1
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.11.1
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.11.1
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.11.1
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.11.1
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.11.1
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.11.1
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.11.1
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.11.1
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.11.1
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-client	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-common	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-examples	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-it	1.2.0-cdh5.11.1

Project	groupId	artifactId	version
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-rest	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-server	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-shell	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-spark	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.11.1
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.11.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.11.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.11.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.11.1
	org.apache.pig	piggybank	0.12.0-cdh5.11.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.11.1
	org.apache.pig	pigunit	0.12.0-cdh5.11.1
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.11.1
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.11.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.11.1
	org.apache.flume	flume-ng-core	1.7.0-cdh5.11.1
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.11.1
	org.apache.flume	flume-ng-node	1.7.0-cdh5.11.1
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.11.1
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.11.1
	org.apache.flume	flume-tools	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-channels	flume-file-channel	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-sources	flume-avro-source	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.11.1
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.11.1

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.11. 1
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.11. 1
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.11. 1
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.11. 1
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.11. 1
	org.apache. flume. flume-ng-sources	flume-jms-source	1.7.0-cdh5.11. 1
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.7.0-cdh5.11. 1
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.7.0-cdh5.11. 1
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.7.0-cdh5.11. 1
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.7.0-cdh5.11. 1
	org.apache. flume. flume-shared	flume-shared-kafka-test	1.7.0-cdh5.11. 1
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.11. 1
	org.apache. oozie	oozie-core	4.1.0-cdh5.11. 1
	org.apache. oozie	oozie-examples	4.1.0-cdh5.11. 1
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.11. 1.oozie-4.1. 0-cdh5.11.1
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.11.

Version and Download Information

Project	groupId	artifactId	version
			1.oozie-4.1.0-cdh5.11.1
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.11.1.oozie-4.1.0-cdh5.11.1
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.11.1.oozie-4.1.0-cdh5.11.1
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.11.1.oozie-4.1.0-cdh5.11.1
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.11.1.oozie-4.1.0-cdh5.11.1
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.11.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.11.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.11.1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.11.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.11.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.11.1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.11.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.11.1
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.11.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.11.1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.11.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.11.1
	org.apache.mahout	mahout-core	0.9-cdh5.11.1
	org.apache.mahout	mahout-examples	0.9-cdh5.11.1

Project	groupId	artifactId	version
	org.apache.mahout	mahout-integration	0.9-cdh5.11.1
	org.apache.mahout	mahout-math	0.9-cdh5.11.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.11.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.11.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.11.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.11.1
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.11.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.11.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.11.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.11.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.11.1
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.11.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.11.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.11.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.11.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.11.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.11.1
	org.apache.sqoop	test	1.99.5-cdh5.11.1
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.11.1
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.11.1
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.11.1
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.11.1
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.11.1

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.11.1
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.11.1
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.11.1
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.11.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.11.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.11.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.11.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.11.1
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.11.1
	com.twitter	parquet-cascading	1.5.0-cdh5.11.1
	com.twitter	parquet-column	1.5.0-cdh5.11.1
	com.twitter	parquet-common	1.5.0-cdh5.11.1
	com.twitter	parquet-encoding	1.5.0-cdh5.11.1
	com.twitter	parquet-format	2.1.0-cdh5.11.1
	com.twitter	parquet-generator	1.5.0-cdh5.11.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.11.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.11.1
	com.twitter	parquet-jackson	1.5.0-cdh5.11.1

Project	groupId	artifactId	version
	com.twitter	parquet-pig	1.5.0-cdh5.11.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.11.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.11.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.11.1
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.11.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.11.1
	com.twitter	parquet-thrift	1.5.0-cdh5.11.1
	com.twitter	parquet-tools	1.5.0-cdh5.11.1
Llama	com.cloudera.llama	llama	1.0.0-cdh5.11.1
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.11.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.11.1
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.11.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.11.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.11.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.11.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.11.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.11.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.11.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.11.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.11.1

Project	groupId	artifactId	version
	org.apache.crunch	crunch-test	0.11.0-cdh5.11.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.11.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.11.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.11.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.11.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.11.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.11.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.11.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.11.1
	org.apache.avro	avro-tools	1.7.6-cdh5.11.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.11.1
	org.apache.avro	trevni-core	1.7.6-cdh5.11.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.11.1
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.11.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.11.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.11.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.11.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.11.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.11.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.11.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.11.1

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.11.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.11.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.11.1
	org.kitesdk	kite-tools	1.0.0-cdh5.11.1

Project	groupId	artifactId	version
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-core	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-join	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.11.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.11.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.11.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.11.1
	org.apache.solr	solr-cell	4.10.3-cdh5.11.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.11.1
	org.apache.solr	solr-core	4.10.3-cdh5.11.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.11.1
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.11.1
	org.apache.solr	solr-langid	4.10.3-cdh5.11.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.11.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.11.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.11.1
	org.apache.solr	solr-uima	4.10.3-cdh5.11.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.11.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.11.1
	com.cloudera.search	search-mr	1.0.0-cdh5.11.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.11.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.11.1

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-common	1.5-cdh5.11.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.11.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.11.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.11.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.11.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.11.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.11.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.11.1
	com.ngdata	hbase-sep-api	1.5-cdh5.11.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.11.1
	com.ngdata	hbase-sep-impl	1.5-hbase1. 2-cdh5.11.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.11.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.11.1

CDH 5.11.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-azure-datalake	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.11.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.11.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.11.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.11.0
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.11.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.11.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.11.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.11.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.11.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.11.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.11.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.11.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.11.0
	org.apache.hive	hive-ant	1.1.0-cdh5.11.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.11.0
	org.apache.hive	hive-cli	1.1.0-cdh5.11.0
	org.apache.hive	hive-common	1.1.0-cdh5.11.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.11.0
	org.apache.hive	hive-exec	1.1.0-cdh5.11.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.11.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.11.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.11.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.11.0
	org.apache.hive	hive-serde	1.1.0-cdh5.11.0
	org.apache.hive	hive-service	1.1.0-cdh5.11.0
	org.apache.hive	hive-shims	1.1.0-cdh5.11.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.11.0
	org.apache.hive	spark-client	1.1.0-cdh5.11.0
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.11.0
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.11.0
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.11.0
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.11.0

Project	groupId	artifactId	version
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.11.0
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.11.0
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.11.0
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.11.0
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.11.0
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-client	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-common	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-examples	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-it	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-rest	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-rsgroup	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-server	1.2.0-cdh5.11.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-shell	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-spark	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.11.0
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.11.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.11.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.11.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.11.0
	org.apache.pig	piggybank	0.12.0-cdh5.11.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.11.0
	org.apache.pig	pigunit	0.12.0-cdh5.11.0
Apache Flume 1.x	org.apache.flume	flume-checkstyle	1.7.0-cdh5.11.0
	org.apache.flume	flume-ng-auth	1.7.0-cdh5.11.0
	org.apache.flume	flume-ng-configuration	1.7.0-cdh5.11.0
	org.apache.flume	flume-ng-core	1.7.0-cdh5.11.0
	org.apache.flume	flume-ng-embedded-agent	1.7.0-cdh5.11.0
	org.apache.flume	flume-ng-node	1.7.0-cdh5.11.0
	org.apache.flume	flume-ng-sdk	1.7.0-cdh5.11.0
	org.apache.flume	flume-ng-tests	1.7.0-cdh5.11.0
	org.apache.flume	flume-tools	1.7.0-cdh5.11.0

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-channels	flume-file-channel	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-channels	flume-jdbc-channel	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-channels	flume-kafka-channel	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-channels	flume-spillable-memory-channel	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-clients	flume-ng-log4jappender	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sources	flume-avro-source	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.7.0-cdh5.11. 0
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.7.0-cdh5.11. 0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-jms-source	1.7.0-cdh5.11.0
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.7.0-cdh5.11.0
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.7.0-cdh5.11.0
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.7.0-cdh5.11.0
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.7.0-cdh5.11.0
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.7.0-cdh5.11.0
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-core	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-examples	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.11.0.oozie-4.1.0-cdh5.11.0
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.11.0.oozie-4.1.0-cdh5.11.0
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.11.0.oozie-4.1.0-cdh5.11.0
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.11.0.oozie-4.1.0-cdh5.11.0
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.11.0.oozie-4.1.0-cdh5.11.0
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.11.0.oozie-4.1.0-cdh5.11.0

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.11.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.11.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.11.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.11.0
	org.apache.mahout	mahout-core	0.9-cdh5.11.0
	org.apache.mahout	mahout-examples	0.9-cdh5.11.0
	org.apache.mahout	mahout-integration	0.9-cdh5.11.0
	org.apache.mahout	mahout-math	0.9-cdh5.11.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.11.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.11.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cli	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.11.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.11.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.11.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.11.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.11.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.11.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.11.0

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.11.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.11.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.11.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.11.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.11.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.11.0
	org.apache.sqoop	test	1.99.5-cdh5.11.0
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.11.0
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.11.0
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.11.0
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.11.0
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.11.0
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.11.0
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.11.0
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.11.0
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.11.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.11.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.11.0

Project	groupId	artifactId	version
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.11.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.11.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.11.0
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.11.0
	com.twitter	parquet-cascading	1.5.0-cdh5.11.0
	com.twitter	parquet-column	1.5.0-cdh5.11.0
	com.twitter	parquet-common	1.5.0-cdh5.11.0
	com.twitter	parquet-encoding	1.5.0-cdh5.11.0
	com.twitter	parquet-format	2.1.0-cdh5.11.0
	com.twitter	parquet-generator	1.5.0-cdh5.11.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.11.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.11.0
	com.twitter	parquet-jackson	1.5.0-cdh5.11.0
	com.twitter	parquet-pig	1.5.0-cdh5.11.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.11.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.11.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.11.0
	com.twitter	parquet-scooage_2.10	1.5.0-cdh5.11.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.11.0
	com.twitter	parquet-thrift	1.5.0-cdh5.11.0

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-tools	1.5.0-cdh5.11.0
Llama	com.cloudera.llama	llama	1.0.0-cdh5.11.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.11.0

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.11.0
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.11.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.11.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.11.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.11.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.11.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.11.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.11.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.11.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.11.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.11.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.11.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.11.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.11.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.11.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.11.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.11.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.11.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.11.0
	org.apache.avro	avro-tools	1.7.6-cdh5.11.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.11.0
	org.apache.avro	trevni-core	1.7.6-cdh5.11.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.11.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.11.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.11.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.11.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.11.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.11.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.11.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.11.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.11.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.11.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.11.0

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.11.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.11.0
	org.kitesdk	kite-tools	1.0.0-cdh5.11.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.11.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.11.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.11.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.11.0

Project	groupId	artifactId	version
	org.apache.solr	solr-cell	4.10.3-cdh5.11.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.11.0
	org.apache.solr	solr-core	4.10.3-cdh5.11.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.11.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.11.0
	org.apache.solr	solr-langid	4.10.3-cdh5.11.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.11.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.11.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.11.0
	org.apache.solr	solr-uima	4.10.3-cdh5.11.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.11.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.11.0
	com.cloudera.search	search-mr	1.0.0-cdh5.11.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.11.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.11.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.11.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.11.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.11.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.11.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.11.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.11.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.11.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.11.0
	com.ngdata	hbase-sep-api	1.5-cdh5.11.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.11.0
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.11.0

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.11.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.11.0

Maven Artifacts for CDH 5.10.x Releases CDH 5.10.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.10.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.10.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.10.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.10.2
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.10.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.10.2
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.10.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.10.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.10.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.10.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.10.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.10.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.10.2
	org.apache.hive	hive-ant	1.1.0-cdh5.10.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.10.2
	org.apache.hive	hive-cli	1.1.0-cdh5.10.2
	org.apache.hive	hive-common	1.1.0-cdh5.10.2

Project	groupId	artifactId	version
	org.apache.hive	hive-contrib	1.1.0-cdh5.10.2
	org.apache.hive	hive-exec	1.1.0-cdh5.10.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.10.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.10.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.10.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.10.2
	org.apache.hive	hive-serde	1.1.0-cdh5.10.2
	org.apache.hive	hive-service	1.1.0-cdh5.10.2
	org.apache.hive	hive-shims	1.1.0-cdh5.10.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.10.2
	org.apache.hive	spark-client	1.1.0-cdh5.10.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.10.2
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.10.2
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.10.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.10.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.10.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.10.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.10.2
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.10.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.10.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.10.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-it	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-rest	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-server	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.10.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.10.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.10.2

Project	groupId	artifactId	version
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.10.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.10.2
	org.apache.pig	piggybank	0.12.0-cdh5.10.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.10.2
	org.apache.pig	pigunit	0.12.0-cdh5.10.2
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.10.2
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.10.2
	org.apache.flume	flume-ng-core	1.6.0-cdh5.10.2
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.10.2
	org.apache.flume	flume-ng-node	1.6.0-cdh5.10.2
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.10.2
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.10.2
	org.apache.flume	flume-tools	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.10.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sources	flume-jms-source	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.10.2
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.10.2
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.6.0-cdh5.10.2

Project	groupId	artifactId	version
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-core	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-examples	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.10.2.oozie-4.1.0-cdh5.10.2
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.10.2.oozie-4.1.0-cdh5.10.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.10.2.oozie-4.1.0-cdh5.10.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.10.2.oozie-4.1.0-cdh5.10.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.10.2.oozie-4.1.0-cdh5.10.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.10.2.oozie-4.1.0-cdh5.10.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.10.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.10.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.10.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.10.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.10.2
	org.apache.mahout	mahout-core	0.9-cdh5.10.2
	org.apache.mahout	mahout-examples	0.9-cdh5.10.2
	org.apache.mahout	mahout-integration	0.9-cdh5.10.2
	org.apache.mahout	mahout-math	0.9-cdh5.10.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.10.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.10.2

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.10.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.10.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.10.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.10.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.10.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.10.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.10.2
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.10.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.10.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.10.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.10.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.10.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.10.2
	org.apache.sqoop	test	1.99.5-cdh5.10.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.10.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.10.2
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.10.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.10.2
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.10.2
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.10.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.10.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.10.2
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.10.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.10.2

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.10.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.10.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.10.2
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.10.2
	com.twitter	parquet-cascading	1.5.0-cdh5.10.2
	com.twitter	parquet-column	1.5.0-cdh5.10.2
	com.twitter	parquet-common	1.5.0-cdh5.10.2

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-encoding	1.5.0-cdh5.10.2
	com.twitter	parquet-format	2.1.0-cdh5.10.2
	com.twitter	parquet-generator	1.5.0-cdh5.10.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.10.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.10.2
	com.twitter	parquet-jackson	1.5.0-cdh5.10.2
	com.twitter	parquet-pig	1.5.0-cdh5.10.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.10.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.10.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.10.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.10.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.10.2
	com.twitter	parquet-thrift	1.5.0-cdh5.10.2
	com.twitter	parquet-tools	1.5.0-cdh5.10.2
Llama	com.cloudera.llama	llama	1.0.0-cdh5.10.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.10.2
	com.databricks	spark-avro_2.10	1.1.0-cdh5.10.2
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.10.2

Project	groupId	artifactId	version
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.10.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.10.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.10.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.10.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.10.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.10.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.10.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.10.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.10.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.10.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.10.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.10.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.10.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.10.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.10.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.10.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.10.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.10.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.10.2
	org.apache.avro	avro-tools	1.7.6-cdh5.10.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.10.2
	org.apache.avro	trevni-core	1.7.6-cdh5.10.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.10.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.10.2

Project	groupId	artifactId	version
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.10.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.10.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.10.2
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.10.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.10.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.10.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.10.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.10.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.10.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.10.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.10.2
	org.kitesdk	kite-tools	1.0.0-cdh5.10.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.10.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-facet	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.10.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.10.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.10.2
	org.apache.solr	solr-cell	4.10.3-cdh5.10.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.10.2
	org.apache.solr	solr-core	4.10.3-cdh5.10.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.10.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.10.2
	org.apache.solr	solr-langid	4.10.3-cdh5.10.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.10.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-solrj	4.10.3-cdh5.10.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.10.2
	org.apache.solr	solr-uima	4.10.3-cdh5.10.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.10.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.10.2
	com.cloudera.search	search-mr	1.0.0-cdh5.10.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.10.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.10.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.10.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.10.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.10.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.10.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.10.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.10.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.10.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.10.2
	com.ngdata	hbase-sep-api	1.5-cdh5.10.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.10.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.10.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.10.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.10.2

CDH 5.10.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.10.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.10.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.10.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.10.1
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.10.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.10.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.10.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.10.1
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.10.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.10.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.10.1
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.10.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.10.1
	org.apache.hive	hive-ant	1.1.0-cdh5.10.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.10.1
	org.apache.hive	hive-cli	1.1.0-cdh5.10.1
	org.apache.hive	hive-common	1.1.0-cdh5.10.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.10.1
	org.apache.hive	hive-exec	1.1.0-cdh5.10.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.10.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.10.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.10.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.10.1
	org.apache.hive	hive-serde	1.1.0-cdh5.10.1
	org.apache.hive	hive-service	1.1.0-cdh5.10.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-shims	1.1.0-cdh5.10.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.10.1
	org.apache.hive	spark-client	1.1.0-cdh5.10.1
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.10.1
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.10.1
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.10.1
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.10.1
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.10.1
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.10.1
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.10.1
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.10.1
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.10.1
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-client	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-common	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-examples	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-it	1.2.0-cdh5.10.1

Project	groupId	artifactId	version
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-rest	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-server	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-shell	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-spark	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.10.1
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.10.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.10.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.10.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.10.1
	org.apache.pig	piggybank	0.12.0-cdh5.10.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.10.1
	org.apache.pig	pigunit	0.12.0-cdh5.10.1
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.10.1
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.10.1
	org.apache.flume	flume-ng-core	1.6.0-cdh5.10.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.10.1
	org.apache.flume	flume-ng-node	1.6.0-cdh5.10.1
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.10.1
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.10.1
	org.apache.flume	flume-tools	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.10.1
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.10.1

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.10. 1
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.10. 1
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.10. 1
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.10. 1
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.10. 1
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.10. 1
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.10. 1
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.10. 1
	org.apache. flume. flume-shared	flume-shared-kafka-test	1.6.0-cdh5.10. 1
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.10. 1
	org.apache. oozie	oozie-core	4.1.0-cdh5.10. 1
	org.apache. oozie	oozie-examples	4.1.0-cdh5.10. 1
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.10. 1.oozie-4.1. 0-cdh5.10.1
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.10. 1.oozie-4.1. 0-cdh5.10.1
	org.apache. oozie	oozie-hadoop-test	2.6. 0-mr1-cdh5.10. 1.oozie-4.1. 0-cdh5.10.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.10.1.oozie-4.1.0-cdh5.10.1
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.10.1.oozie-4.1.0-cdh5.10.1
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.10.1.oozie-4.1.0-cdh5.10.1
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.10.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.10.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.10.1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.10.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.10.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.10.1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.10.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.10.1
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.10.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.10.1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.10.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.10.1
	org.apache.mahout	mahout-core	0.9-cdh5.10.1
	org.apache.mahout	mahout-examples	0.9-cdh5.10.1
	org.apache.mahout	mahout-integration	0.9-cdh5.10.1
	org.apache.mahout	mahout-math	0.9-cdh5.10.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.10.1

Project	groupId	artifactId	version
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.10.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.10.1
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.10.1

Version and Download Information

Project	groupId	artifactId	version
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.10.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.10.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.10.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.10.1
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.10.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.10.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.10.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.10.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.10.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.10.1
	org.apache.sqoop	test	1.99.5-cdh5.10.1
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.10.1
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.10.1
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.10.1
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.10.1
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.10.1
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.10.1
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.10.1

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.10.1
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.10.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.10.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.10.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.10.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.10.1
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.10.1
	com.twitter	parquet-cascading	1.5.0-cdh5.10.1
	com.twitter	parquet-column	1.5.0-cdh5.10.1
	com.twitter	parquet-common	1.5.0-cdh5.10.1
	com.twitter	parquet-encoding	1.5.0-cdh5.10.1
	com.twitter	parquet-format	2.1.0-cdh5.10.1
	com.twitter	parquet-generator	1.5.0-cdh5.10.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.10.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.10.1
	com.twitter	parquet-jackson	1.5.0-cdh5.10.1
	com.twitter	parquet-pig	1.5.0-cdh5.10.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.10.1

Project	groupId	artifactId	version
	com.twitter	parquet-protobuf	1.5.0-cdh5.10.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.10.1
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.10.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.10.1
	com.twitter	parquet-thrift	1.5.0-cdh5.10.1
	com.twitter	parquet-tools	1.5.0-cdh5.10.1
Llama	com.cloudera.llama	llama	1.0.0-cdh5.10.1
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.10.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.10.1
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.10.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.10.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.10.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.10.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.10.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.10.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.10.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.10.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.10.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.10.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.10.1

Project	groupId	artifactId	version
	org.apache.avro	avro-compiler	1.7.6-cdh5.10.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.10.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.10.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.10.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.10.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.10.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.10.1
	org.apache.avro	avro-tools	1.7.6-cdh5.10.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.10.1
	org.apache.avro	trevni-core	1.7.6-cdh5.10.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.10.1
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.10.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.10.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.10.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.10.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.10.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.10.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.10.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.10.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.10.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.10.1

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.10.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.10.1
	org.kitesdk	kite-tools	1.0.0-cdh5.10.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.10.1

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-core	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-join	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.10.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.10.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.10.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.10.1
	org.apache.solr	solr-cell	4.10.3-cdh5.10.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.10.1
	org.apache.solr	solr-core	4.10.3-cdh5.10.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.10.1
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.10.1
	org.apache.solr	solr-langid	4.10.3-cdh5.10.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.10.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.10.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.10.1
	org.apache.solr	solr-uima	4.10.3-cdh5.10.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.10.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.10.1
	com.cloudera.search	search-mr	1.0.0-cdh5.10.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.10.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.10.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.10.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.10.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.10.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.10.1

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-model	1.5-cdh5.10.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.10.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.10.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.10.1
	com.ngdata	hbase-sep-api	1.5-cdh5.10.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.10.1
	com.ngdata	hbase-sep-impl	1.5-hbase1. 2-cdh5.10.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.10.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.10.1

CDH 5.10.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.10.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.10.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.10.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.10.0
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.10.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.10.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.10.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.10.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.10.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.10.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.10.0
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.10.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.10.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-ant	1.1.0-cdh5.10.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.10.0
	org.apache.hive	hive-cli	1.1.0-cdh5.10.0
	org.apache.hive	hive-common	1.1.0-cdh5.10.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.10.0
	org.apache.hive	hive-exec	1.1.0-cdh5.10.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.10.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.10.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.10.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.10.0
	org.apache.hive	hive-serde	1.1.0-cdh5.10.0
	org.apache.hive	hive-service	1.1.0-cdh5.10.0
	org.apache.hive	hive-shims	1.1.0-cdh5.10.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.10.0
	org.apache.hive	spark-client	1.1.0-cdh5.10.0
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.10.0
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.10.0
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.10.0
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.10.0
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.10.0
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.10.0

Project	groupId	artifactId	version
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.10.0
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.10.0
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.10.0
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-client	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-common	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-examples	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-it	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-rest	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-server	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-shell	1.2.0-cdh5.10.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-spark	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.10.0
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.10.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.10.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.10.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.10.0
	org.apache.pig	piggybank	0.12.0-cdh5.10.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.10.0
	org.apache.pig	pigunit	0.12.0-cdh5.10.0
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.10.0
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.10.0
	org.apache.flume	flume-ng-core	1.6.0-cdh5.10.0
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.10.0
	org.apache.flume	flume-ng-node	1.6.0-cdh5.10.0
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.10.0
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.10.0
	org.apache.flume	flume-tools	1.6.0-cdh5.10.0
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.10.0
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.10.0
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.10.0

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.10. 0
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.10. 0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.10.0
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.10.0
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-core	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-examples	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.10.0.oozie-4.1.0-cdh5.10.0
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.10.0.oozie-4.1.0-cdh5.10.0
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.10.0.oozie-4.1.0-cdh5.10.0
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.10.0.oozie-4.1.0-cdh5.10.0
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.10.0.oozie-4.1.0-cdh5.10.0
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.10.0.oozie-4.1.0-cdh5.10.0
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.10.0

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.10.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.10.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.10.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.10.0
	org.apache.mahout	mahout-core	0.9-cdh5.10.0
	org.apache.mahout	mahout-examples	0.9-cdh5.10.0
	org.apache.mahout	mahout-integration	0.9-cdh5.10.0
	org.apache.mahout	mahout-math	0.9-cdh5.10.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.10.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.10.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.10.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.10.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.10.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.10.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.10.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.10.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.10.0
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.10.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.10.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.10.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.10.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.10.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.10.0

Project	groupId	artifactId	version
	org.apache.sqoop	test	1.99.5-cdh5.10.0
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.10.0
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.10.0
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.10.0
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.10.0
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.10.0
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.10.0
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.10.0
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.10.0
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.10.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.10.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.10.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.10.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.10.0
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.10.0
	com.twitter	parquet-cascading	1.5.0-cdh5.10.0

Project	groupId	artifactId	version
	com.twitter	parquet-column	1.5.0-cdh5.10.0
	com.twitter	parquet-common	1.5.0-cdh5.10.0
	com.twitter	parquet-encoding	1.5.0-cdh5.10.0
	com.twitter	parquet-format	2.1.0-cdh5.10.0
	com.twitter	parquet-generator	1.5.0-cdh5.10.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.10.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.10.0
	com.twitter	parquet-jackson	1.5.0-cdh5.10.0
	com.twitter	parquet-pig	1.5.0-cdh5.10.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.10.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.10.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.10.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.10.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.10.0
	com.twitter	parquet-thrift	1.5.0-cdh5.10.0
	com.twitter	parquet-tools	1.5.0-cdh5.10.0
Llama	com.cloudera.llama	llama	1.0.0-cdh5.10.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.10.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.10.0
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.10.0

Project	groupId	artifactId	version
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.10.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.10.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.10.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.10.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.10.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.10.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.10.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.10.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.10.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.10.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.10.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.10.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.10.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.10.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.10.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.10.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.10.0
	org.apache.avro	avro-tools	1.7.6-cdh5.10.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.10.0
	org.apache.avro	trevni-core	1.7.6-cdh5.10.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.10.0

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.10.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.10.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.10.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.10.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.10.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.10.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.10.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.10.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.10.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.10.0

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.10.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.10.0
	org.kitesdk	kite-tools	1.0.0-cdh5.10.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.10.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.10.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.10.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.10.0
	org.apache.solr	solr-cell	4.10.3-cdh5.10.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.10.0
	org.apache.solr	solr-core	4.10.3-cdh5.10.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.10.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.10.0
	org.apache.solr	solr-langid	4.10.3-cdh5.10.0

Project	groupId	artifactId	version
	org.apache.solr	solr-security-util	4.10.3-cdh5.10.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.10.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.10.0
	org.apache.solr	solr-uima	4.10.3-cdh5.10.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.10.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.10.0
	com.cloudera.search	search-mr	1.0.0-cdh5.10.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.10.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.10.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.10.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.10.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.10.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.10.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.10.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.10.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.10.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.10.0
	com.ngdata	hbase-sep-api	1.5-cdh5.10.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.10.0
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.10.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.10.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.10.0

Maven Artifacts for CDH 5.9.x Releases CDH 5.9.3

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.9.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.9.3

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.9.3
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.9.3
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.9.3

Version and Download Information

Project	groupId	artifactId	version
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.9.3
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.9.3
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.9.3
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.9.3
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.9.3
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.9.3
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.9.3
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.9.3
	org.apache.hive	hive-ant	1.1.0-cdh5.9.3
	org.apache.hive	hive-beeline	1.1.0-cdh5.9.3
	org.apache.hive	hive-cli	1.1.0-cdh5.9.3
	org.apache.hive	hive-common	1.1.0-cdh5.9.3
	org.apache.hive	hive-contrib	1.1.0-cdh5.9.3
	org.apache.hive	hive-exec	1.1.0-cdh5.9.3
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.9.3
	org.apache.hive	hive-hwi	1.1.0-cdh5.9.3
	org.apache.hive	hive-jdbc	1.1.0-cdh5.9.3
	org.apache.hive	hive-metastore	1.1.0-cdh5.9.3

Project	groupId	artifactId	version
	org.apache.hive	hive-serde	1.1.0-cdh5.9.3
	org.apache.hive	hive-service	1.1.0-cdh5.9.3
	org.apache.hive	hive-shims	1.1.0-cdh5.9.3
	org.apache.hive	hive-testutils	1.1.0-cdh5.9.3
	org.apache.hive	spark-client	1.1.0-cdh5.9.3
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.9.3
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.9.3
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.9.3
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.9.3
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.9.3
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.9.3
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.9.3
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.9.3
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.9.3
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-client	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-common	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-examples	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.9.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-it	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-rest	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-server	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-shell	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-spark	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.9.3
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.9.3
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.9.3
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.9.3
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.9.3
	org.apache.pig	piggybank	0.12.0-cdh5.9.3
	org.apache.pig	pigsmoke	0.12.0-cdh5.9.3
	org.apache.pig	pigunit	0.12.0-cdh5.9.3
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.9.3

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.9.3
	org.apache.flume	flume-ng-core	1.6.0-cdh5.9.3
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.9.3
	org.apache.flume	flume-ng-node	1.6.0-cdh5.9.3
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.9.3
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.9.3
	org.apache.flume	flume-tools	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.9.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sources	flume-jms-source	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.9.3
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.9.3
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.6.0-cdh5.9.3
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-core	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-examples	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.9.3.oozie-4.1.0-cdh5.9.3
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.9.

Project	groupId	artifactId	version
			3.oozie-4.1.0-cdh5.9.3
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.9.3.oozie-4.1.0-cdh5.9.3
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.9.3.oozie-4.1.0-cdh5.9.3
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.9.3.oozie-4.1.0-cdh5.9.3
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.9.3.oozie-4.1.0-cdh5.9.3
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.9.3
	org.apache.oozie	oozie-tools	4.1.0-cdh5.9.3
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.9.3
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.9.3
	org.apache.mahout	mahout-core	0.9-cdh5.9.3
	org.apache.mahout	mahout-examples	0.9-cdh5.9.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.mahout	mahout-integration	0.9-cdh5.9.3
	org.apache.mahout	mahout-math	0.9-cdh5.9.3
	org.apache.mahout	mahout-math-scala	0.9-cdh5.9.3
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-chef	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-cli	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-core	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-druid	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-examples	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-hama	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-pig	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-solr	0.9.0-cdh5.9.3

Project	groupId	artifactId	version
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.9.3
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.9.3
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.9.3
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.9.3
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.9.3
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.9.3
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.9.3
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.9.3
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.9.3
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.9.3
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.9.3
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.9.3
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.9.3
	org.apache.sqoop	test	1.99.5-cdh5.9.3
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.9.3
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.9.3
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.9.3
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.9.3
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.9.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.9.3
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.9.3
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.9.3
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.9.3
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-dist	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.9.3

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.9.3
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.9.3
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.9.3
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.9.3
	com.twitter	parquet-cascading	1.5.0-cdh5.9.3
	com.twitter	parquet-column	1.5.0-cdh5.9.3
	com.twitter	parquet-common	1.5.0-cdh5.9.3
	com.twitter	parquet-encoding	1.5.0-cdh5.9.3
	com.twitter	parquet-format	2.1.0-cdh5.9.3
	com.twitter	parquet-generator	1.5.0-cdh5.9.3
	com.twitter	parquet-hadoop	1.5.0-cdh5.9.3
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.9.3
	com.twitter	parquet-jackson	1.5.0-cdh5.9.3
	com.twitter	parquet-pig	1.5.0-cdh5.9.3
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.9.3
	com.twitter	parquet-protobuf	1.5.0-cdh5.9.3
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.9.3
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.9.3
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.9.3
	com.twitter	parquet-thrift	1.5.0-cdh5.9.3

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-tools	1.5.0-cdh5.9.3
Llama	com.cloudera. llama	llama	1.0.0-cdh5.9.3
Apache Spark	org.apache. spark	spark-assembly_2.10	1.6.0-cdh5.9.3
	com.databricks	spark-avro_2.10	1.1.0-cdh5.9.3
	org.apache. spark	spark-bagel_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-catalyst_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-core_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-graphx_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-hive_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-launcher_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-mllib_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-network-common_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-network-shuffle_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-network-yarn_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-repl_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-sql_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-streaming-flume_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-streaming-kafka_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.9.3
	org.apache. spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.9.3

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.9.3
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.9.3
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.9.3
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.9.3
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.9.3
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.9.3
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.9.3
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.9.3
	org.apache.crunch	crunch-core	0.11.0-cdh5.9.3
	org.apache.crunch	crunch-examples	0.11.0-cdh5.9.3
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.9.3
	org.apache.crunch	crunch-hive	0.11.0-cdh5.9.3
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.9.3
	org.apache.crunch	crunch-spark	0.11.0-cdh5.9.3
	org.apache.crunch	crunch-test	0.11.0-cdh5.9.3
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.9.3
	org.apache.avro	avro-compiler	1.7.6-cdh5.9.3
	org.apache.avro	avro-ipc	1.7.6-cdh5.9.3
	org.apache.avro	avro-mapred	1.7.6-cdh5.9.3
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.9.3
	org.apache.avro	avro-protobuf	1.7.6-cdh5.9.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.9.3
	org.apache.avro	avro-thrift	1.7.6-cdh5.9.3
	org.apache.avro	avro-tools	1.7.6-cdh5.9.3
	org.apache.avro	trevni-avro	1.7.6-cdh5.9.3
	org.apache.avro	trevni-core	1.7.6-cdh5.9.3
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.9.3
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.9.3
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.9.3
	org.kitesdk	kite-data-hive	1.0.0-cdh5.9.3
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.9.3
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.9.3
	org.kitesdk	kite-data-s3	1.0.0-cdh5.9.3
	org.kitesdk	kite-data-spark	1.0.0-cdh5.9.3
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.9.3
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.9.3
	org.kitesdk	kite-minicluster	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.9.3

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.9.3
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.9.3
	org.kitesdk	kite-tools	1.0.0-cdh5.9.3
Apache Solr	org.apache. lucene	lucene-analyzers-common	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-analyzers-icu	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-analyzers-morfologik	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-analyzers-phonetic	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-analyzers-smartcn	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-analyzers-stempel	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-analyzers-uima	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-benchmark	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-classification	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-codecs	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-core	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-demo	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-expressions	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-facet	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-grouping	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-highlighter	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-join	4.10.3-cdh5.9. 3
	org.apache. lucene	lucene-memory	4.10.3-cdh5.9. 3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-misc	4.10.3-cdh5.9.3
	org.apache.lucene	lucene-queries	4.10.3-cdh5.9.3
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.9.3
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.9.3
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.9.3
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.9.3
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.9.3
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.9.3
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.9.3
	org.apache.solr	solr-cell	4.10.3-cdh5.9.3
	org.apache.solr	solr-clustering	4.10.3-cdh5.9.3
	org.apache.solr	solr-core	4.10.3-cdh5.9.3
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.9.3
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.9.3
	org.apache.solr	solr-langid	4.10.3-cdh5.9.3
	org.apache.solr	solr-security-util	4.10.3-cdh5.9.3
	org.apache.solr	solr-solrj	4.10.3-cdh5.9.3
	org.apache.solr	solr-test-framework	4.10.3-cdh5.9.3
	org.apache.solr	solr-uima	4.10.3-cdh5.9.3
	org.apache.solr	solr-velocity	4.10.3-cdh5.9.3
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.9.3

Project	groupId	artifactId	version
	com.cloudera. search	search-mr	1.0.0-cdh5.9.3
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.9.3
	com.ngdata	hbase-indexer-cli	1.5-cdh5.9.3
	com.ngdata	hbase-indexer-common	1.5-cdh5.9.3
	com.ngdata	hbase-indexer-demo	1.5-cdh5.9.3
	com.ngdata	hbase-indexer-dist	1.5-cdh5.9.3
	com.ngdata	hbase-indexer-engine	1.5-cdh5.9.3
	com.ngdata	hbase-indexer-model	1.5-cdh5.9.3
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.9.3
	com.ngdata	hbase-indexer-mr	1.5-cdh5.9.3
	com.ngdata	hbase-indexer-server	1.5-cdh5.9.3
	com.ngdata	hbase-sep-api	1.5-cdh5.9.3
	com.ngdata	hbase-sep-demo	1.5-cdh5.9.3
	com.ngdata	hbase-sep-impl	1.5-hbase1. 2-cdh5.9.3
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.9.3
	com.ngdata	hbase-sep-tools	1.5-cdh5.9.3

CDH 5.9.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache. hadoop	hadoop-annotations	2.6.0-cdh5.9.2
	org.apache. hadoop	hadoop-ant	2.6.0-cdh5.9.2
	org.apache. hadoop	hadoop-archive-logs	2.6.0-cdh5.9.2
	org.apache. hadoop	hadoop-archives	2.6.0-cdh5.9.2
	org.apache. hadoop	hadoop-assemblies	2.6.0-cdh5.9.2
	org.apache. hadoop	hadoop-auth	2.6.0-cdh5.9.2
	org.apache. hadoop	hadoop-aws	2.6.0-cdh5.9.2
	org.apache. hadoop	hadoop-azure	2.6.0-cdh5.9.2
	org.apache. hadoop	hadoop-build-tools	2.6.0-cdh5.9.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.9.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.9.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.9.2
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.9.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.9.2
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.9.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.9.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.9.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.9.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.9.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.9.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.9.2
	org.apache.hive	hive-ant	1.1.0-cdh5.9.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.9.2
	org.apache.hive	hive-cli	1.1.0-cdh5.9.2
	org.apache.hive	hive-common	1.1.0-cdh5.9.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.9.2
	org.apache.hive	hive-exec	1.1.0-cdh5.9.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.9.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.9.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.9.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.9.2
	org.apache.hive	hive-serde	1.1.0-cdh5.9.2
	org.apache.hive	hive-service	1.1.0-cdh5.9.2
	org.apache.hive	hive-shims	1.1.0-cdh5.9.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.9.2
	org.apache.hive	spark-client	1.1.0-cdh5.9.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.9.2
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.9.2

Project	groupId	artifactId	version
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.9.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.9.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.9.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.9.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.9.2
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.9.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.9.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-it	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-rest	1.2.0-cdh5.9.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-server	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.9.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.9.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.9.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.9.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.9.2
	org.apache.pig	piggybank	0.12.0-cdh5.9.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.9.2
	org.apache.pig	pigunit	0.12.0-cdh5.9.2
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.9.2
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.9.2
	org.apache.flume	flume-ng-core	1.6.0-cdh5.9.2
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.9.2
	org.apache.flume	flume-ng-node	1.6.0-cdh5.9.2
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.9.2
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.9.2
	org.apache.flume	flume-tools	1.6.0-cdh5.9.2

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-channels	flume-file-channel	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.9.2
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.9.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-jms-source	1.6.0-cdh5.9.2
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.6.0-cdh5.9.2
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.9.2
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.9.2
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.9.2
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.6.0-cdh5.9.2
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-core	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-examples	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.9.2.oozie-4.1.0-cdh5.9.2
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.9.2.oozie-4.1.0-cdh5.9.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.9.2.oozie-4.1.0-cdh5.9.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.9.2.oozie-4.1.0-cdh5.9.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.9.2.oozie-4.1.0-cdh5.9.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.9.2.oozie-4.1.0-cdh5.9.2

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.9.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.9.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.9.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.9.2
	org.apache.mahout	mahout-core	0.9-cdh5.9.2
	org.apache.mahout	mahout-examples	0.9-cdh5.9.2
	org.apache.mahout	mahout-integration	0.9-cdh5.9.2
	org.apache.mahout	mahout-math	0.9-cdh5.9.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.9.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.9.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cli	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.9.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.9.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.9.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.9.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.9.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.9.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.9.2

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.9.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.9.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.9.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.9.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.9.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.9.2
	org.apache.sqoop	test	1.99.5-cdh5.9.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.9.2
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.9.2
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.9.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.9.2
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.9.2
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.9.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.9.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.9.2
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.9.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.9.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.9.2

Project	groupId	artifactId	version
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.9.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.9.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.9.2
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.9.2
	com.twitter	parquet-cascading	1.5.0-cdh5.9.2
	com.twitter	parquet-column	1.5.0-cdh5.9.2
	com.twitter	parquet-common	1.5.0-cdh5.9.2
	com.twitter	parquet-encoding	1.5.0-cdh5.9.2
	com.twitter	parquet-format	2.1.0-cdh5.9.2
	com.twitter	parquet-generator	1.5.0-cdh5.9.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.9.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.9.2
	com.twitter	parquet-jackson	1.5.0-cdh5.9.2
	com.twitter	parquet-pig	1.5.0-cdh5.9.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.9.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.9.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.9.2
	com.twitter	parquet-scooage_2.10	1.5.0-cdh5.9.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.9.2
	com.twitter	parquet-thrift	1.5.0-cdh5.9.2
	com.twitter	parquet-tools	1.5.0-cdh5.9.2
Llama	com.cloudera.llama	llama	1.0.0-cdh5.9.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.9.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.9.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.9.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.9.2

Project	groupId	artifactId	version
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.9.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.9.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.9.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.9.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.9.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.9.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.9.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.9.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.9.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.9.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.9.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.9.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.9.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.9.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.9.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.9.2
	org.apache.avro	avro-tools	1.7.6-cdh5.9.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.9.2
	org.apache.avro	trevni-core	1.7.6-cdh5.9.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.9.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.9.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.9.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.9.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.9.2
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.9.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.9.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.9.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.9.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.9.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.9.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.9.2
	org.kitesdk	kite-tools	1.0.0-cdh5.9.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.9.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.9.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.9.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.9.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.9.2
	org.apache.solr	solr-cell	4.10.3-cdh5.9.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.9.2
	org.apache.solr	solr-core	4.10.3-cdh5.9.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.9.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.9.2
	org.apache.solr	solr-langid	4.10.3-cdh5.9.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.9.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.9.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.9.2
	org.apache.solr	solr-uima	4.10.3-cdh5.9.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.9.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.9.2
	com.cloudera.search	search-mr	1.0.0-cdh5.9.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.9.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.9.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.9.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.9.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.9.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.9.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.9.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.9.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.9.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.9.2
	com.ngdata	hbase-sep-api	1.5-cdh5.9.2

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-demo	1.5-cdh5.9.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.9.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.9.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.9.2

CDH 5.9.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.9.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.9.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.9.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.9.1
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.9.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.9.1
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.9.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.9.1
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.9.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.9.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.9.1
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.9.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.9.1
	org.apache.hive	hive-ant	1.1.0-cdh5.9.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.9.1
	org.apache.hive	hive-cli	1.1.0-cdh5.9.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-common	1.1.0-cdh5.9.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.9.1
	org.apache.hive	hive-exec	1.1.0-cdh5.9.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.9.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.9.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.9.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.9.1
	org.apache.hive	hive-serde	1.1.0-cdh5.9.1
	org.apache.hive	hive-service	1.1.0-cdh5.9.1
	org.apache.hive	hive-shims	1.1.0-cdh5.9.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.9.1
	org.apache.hive	spark-client	1.1.0-cdh5.9.1
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.9.1
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.9.1
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.9.1
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.9.1
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.9.1
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.9.1
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.9.1
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.9.1
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.9.1

Project	groupId	artifactId	version
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-client	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-common	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-examples	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-it	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-rest	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-server	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-shell	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-spark	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.9.1
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.9.1

Version and Download Information

Project	groupId	artifactId	version
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.9.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.9.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.9.1
	org.apache.pig	piggybank	0.12.0-cdh5.9.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.9.1
	org.apache.pig	pigunit	0.12.0-cdh5.9.1
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.9.1
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.9.1
	org.apache.flume	flume-ng-core	1.6.0-cdh5.9.1
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.9.1
	org.apache.flume	flume-ng-node	1.6.0-cdh5.9.1
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.9.1
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.9.1
	org.apache.flume	flume-tools	1.6.0-cdh5.9.1
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.9.1
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.9.1
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.9.1
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.9.1
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.9.1

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.9.1
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.9.1

Version and Download Information

Project	groupId	artifactId	version
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-core	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-examples	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.9.1.oozie-4.1.0-cdh5.9.1
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.9.1.oozie-4.1.0-cdh5.9.1
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.9.1.oozie-4.1.0-cdh5.9.1
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.9.1.oozie-4.1.0-cdh5.9.1
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.9.1.oozie-4.1.0-cdh5.9.1
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.9.1.oozie-4.1.0-cdh5.9.1
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.9.1

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.9.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.9.1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.9.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.9.1
	org.apache.mahout	mahout-core	0.9-cdh5.9.1
	org.apache.mahout	mahout-examples	0.9-cdh5.9.1
	org.apache.mahout	mahout-integration	0.9-cdh5.9.1
	org.apache.mahout	mahout-math	0.9-cdh5.9.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.9.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.9.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.9.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.9.1
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.9.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.9.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.9.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.9.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.9.1
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.9.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.9.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.9.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.9.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.9.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.9.1
	org.apache.sqoop	test	1.99.5-cdh5.9.1
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.9.1

Project	groupId	artifactId	version
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.9.1
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.9.1
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.9.1
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.9.1
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.9.1
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.9.1
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.9.1
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.9.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.9.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.9.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.9.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.9.1
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.9.1
	com.twitter	parquet-cascading	1.5.0-cdh5.9.1
	com.twitter	parquet-column	1.5.0-cdh5.9.1
	com.twitter	parquet-common	1.5.0-cdh5.9.1
	com.twitter	parquet-encoding	1.5.0-cdh5.9.1
	com.twitter	parquet-format	2.1.0-cdh5.9.1
	com.twitter	parquet-generator	1.5.0-cdh5.9.1

Project	groupId	artifactId	version
	com.twitter	parquet-hadoop	1.5.0-cdh5.9.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.9.1
	com.twitter	parquet-jackson	1.5.0-cdh5.9.1
	com.twitter	parquet-pig	1.5.0-cdh5.9.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.9.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.9.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.9.1
	com.twitter	parquet-scooage_2.10	1.5.0-cdh5.9.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.9.1
	com.twitter	parquet-thrift	1.5.0-cdh5.9.1
	com.twitter	parquet-tools	1.5.0-cdh5.9.1
Llama	com.cloudera. llama	llama	1.0.0-cdh5.9.1
Apache Spark	org.apache. spark	spark-assembly_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-bagel_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-catalyst_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-core_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-graphx_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-hive_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-launcher_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-mllib_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-network-common_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-network-shuffle_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-network-yarn_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-repl_2.10	1.6.0-cdh5.9.1
	org.apache. spark	spark-sql_2.10	1.6.0-cdh5.9.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.9.1
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.9.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.9.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.9.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.9.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.9.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.9.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.9.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.9.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.9.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.9.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.9.1

Project	groupId	artifactId	version
	org.apache.avro	avro-compiler	1.7.6-cdh5.9.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.9.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.9.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.9.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.9.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.9.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.9.1
	org.apache.avro	avro-tools	1.7.6-cdh5.9.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.9.1
	org.apache.avro	trevni-core	1.7.6-cdh5.9.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.9.1
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.9.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.9.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.9.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.9.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.9.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.9.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.9.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.9.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.9.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.9.1

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.9.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.9.1
	org.kitesdk	kite-tools	1.0.0-cdh5.9.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-core	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.9.1

Project	groupId	artifactId	version
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-join	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.9.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.9.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.9.1
	org.apache.solr	solr-cell	4.10.3-cdh5.9.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.9.1
	org.apache.solr	solr-core	4.10.3-cdh5.9.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.9.1
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.9.1
	org.apache.solr	solr-langid	4.10.3-cdh5.9.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.9.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.9.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-test-framework	4.10.3-cdh5.9.1
	org.apache.solr	solr-uima	4.10.3-cdh5.9.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.9.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.9.1
	com.cloudera.search	search-mr	1.0.0-cdh5.9.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.9.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.9.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.9.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.9.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.9.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.9.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.9.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.9.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.9.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.9.1
	com.ngdata	hbase-sep-api	1.5-cdh5.9.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.9.1
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.9.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.9.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.9.1

CDH 5.9.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.9.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-build-tools	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.9.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.9.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.9.0
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.9.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.9.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.9.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.9.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.9.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.9.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.9.0
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.9.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.9.0
	org.apache.hive	hive-ant	1.1.0-cdh5.9.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.9.0
	org.apache.hive	hive-cli	1.1.0-cdh5.9.0
	org.apache.hive	hive-common	1.1.0-cdh5.9.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.9.0
	org.apache.hive	hive-exec	1.1.0-cdh5.9.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.9.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.9.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.9.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.9.0
	org.apache.hive	hive-serde	1.1.0-cdh5.9.0
	org.apache.hive	hive-service	1.1.0-cdh5.9.0
	org.apache.hive	hive-shims	1.1.0-cdh5.9.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-testutils	1.1.0-cdh5.9.0
	org.apache.hive	spark-client	1.1.0-cdh5.9.0
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.9.0
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.9.0
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.9.0
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.9.0
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.9.0
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.9.0
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.9.0
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.9.0
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.9.0
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-client	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-common	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-examples	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-it	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.9.0

Project	groupId	artifactId	version
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-rest	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-server	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-shell	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-spark	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.9.0
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.9.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.9.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.9.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.9.0
	org.apache.pig	piggybank	0.12.0-cdh5.9.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.9.0
	org.apache.pig	pigunit	0.12.0-cdh5.9.0
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.9.0
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.9.0
	org.apache.flume	flume-ng-core	1.6.0-cdh5.9.0
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.9.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-node	1.6.0-cdh5.9.0
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.9.0
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.9.0
	org.apache.flume	flume-tools	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.9.0
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.9.0

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.9.0
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.9.0
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.9.0
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.9.0
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.9.0
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.9.0
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.9.0
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.9.0
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.9.0
	org.apache. oozie	oozie-core	4.1.0-cdh5.9.0
	org.apache. oozie	oozie-examples	4.1.0-cdh5.9.0
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.9. 0.oozie-4.1. 0-cdh5.9.0
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.9. 0.oozie-4.1. 0-cdh5.9.0
	org.apache. oozie	oozie-hadoop-test	2.6. 0-mr1-cdh5.9. 0.oozie-4.1. 0-cdh5.9.0
	org.apache. oozie	oozie-hadoop-utils	2.6. 0-mr1-cdh5.9. 0.oozie-4.1. 0-cdh5.9.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.9.0.oozie-4.1.0-cdh5.9.0
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.9.0.oozie-4.1.0-cdh5.9.0
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.9.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.9.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.9.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.9.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.9.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.9.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.9.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.9.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.9.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.9.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.9.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.9.0
	org.apache.mahout	mahout-core	0.9-cdh5.9.0
	org.apache.mahout	mahout-examples	0.9-cdh5.9.0
	org.apache.mahout	mahout-integration	0.9-cdh5.9.0
	org.apache.mahout	mahout-math	0.9-cdh5.9.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.9.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.9.0

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.9.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.9.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.9.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.9.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.9.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.9.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.9.0
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.9.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.9.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.9.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.9.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.9.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.9.0
	org.apache.sqoop	test	1.99.5-cdh5.9.0
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.9.0
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.9.0
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.9.0
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.9.0
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.9.0
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.9.0
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.9.0
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.9.0

Project	groupId	artifactId	version
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.9.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.9.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.9.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.9.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.9.0
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.9.0
	com.twitter	parquet-cascading	1.5.0-cdh5.9.0
	com.twitter	parquet-column	1.5.0-cdh5.9.0
	com.twitter	parquet-common	1.5.0-cdh5.9.0
	com.twitter	parquet-encoding	1.5.0-cdh5.9.0
	com.twitter	parquet-format	2.1.0-cdh5.9.0
	com.twitter	parquet-generator	1.5.0-cdh5.9.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.9.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.9.0
	com.twitter	parquet-jackson	1.5.0-cdh5.9.0
	com.twitter	parquet-pig	1.5.0-cdh5.9.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.9.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.9.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.9.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.9.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.9.0
	com.twitter	parquet-thrift	1.5.0-cdh5.9.0
	com.twitter	parquet-tools	1.5.0-cdh5.9.0
Llama	com.cloudera.llama	llama	1.0.0-cdh5.9.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.9.0

Project	groupId	artifactId	version
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.9.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.9.0
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.9.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.9.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.9.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.9.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.9.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.9.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.9.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.9.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.9.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.9.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.9.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.9.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.9.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.9.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.9.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.9.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.9.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.9.0
	org.apache.avro	avro-tools	1.7.6-cdh5.9.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.9.0

Project	groupId	artifactId	version
	org.apache.avro	trevni-core	1.7.6-cdh5.9.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.9.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.9.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.9.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.9.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.9.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.9.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.9.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.9.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.9.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.9.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.9.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.9.0
	org.kitesdk	kite-tools	1.0.0-cdh5.9.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.9.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.9.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.9.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.9.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.9.0
	org.apache.solr	solr-cell	4.10.3-cdh5.9.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.9.0
	org.apache.solr	solr-core	4.10.3-cdh5.9.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.9.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.9.0
	org.apache.solr	solr-langid	4.10.3-cdh5.9.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.9.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.9.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.9.0
	org.apache.solr	solr-uima	4.10.3-cdh5.9.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.9.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.9.0
	com.cloudera.search	search-mr	1.0.0-cdh5.9.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.9.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.9.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.9.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.9.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.9.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.9.0

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-model	1.5-cdh5.9.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.9.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.9.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.9.0
	com.ngdata	hbase-sep-api	1.5-cdh5.9.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.9.0
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.9.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.9.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.9.0

Maven Artifacts for CDH 5.8.x Releases CDH 5.8.5

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.8.5

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.8.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.8.5
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.8.5
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.8.5
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.8.5
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.8.5
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.8.5
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.8.5
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.8.5
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.8.5
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.8.5
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.8.5
	org.apache.hive	hive-ant	1.1.0-cdh5.8.5

Project	groupId	artifactId	version
	org.apache.hive	hive-beeline	1.1.0-cdh5.8.5
	org.apache.hive	hive-cli	1.1.0-cdh5.8.5
	org.apache.hive	hive-common	1.1.0-cdh5.8.5
	org.apache.hive	hive-contrib	1.1.0-cdh5.8.5
	org.apache.hive	hive-exec	1.1.0-cdh5.8.5
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.8.5
	org.apache.hive	hive-hwi	1.1.0-cdh5.8.5
	org.apache.hive	hive-jdbc	1.1.0-cdh5.8.5
	org.apache.hive	hive-metastore	1.1.0-cdh5.8.5
	org.apache.hive	hive-serde	1.1.0-cdh5.8.5
	org.apache.hive	hive-service	1.1.0-cdh5.8.5
	org.apache.hive	hive-shims	1.1.0-cdh5.8.5
	org.apache.hive	hive-testutils	1.1.0-cdh5.8.5
	org.apache.hive	spark-client	1.1.0-cdh5.8.5
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.8.5
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.8.5
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.8.5
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.8.5
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.8.5
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.8.5
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.8.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.8.5
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.8.5
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-client	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-common	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-examples	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-it	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-rest	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-server	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-shell	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-spark	1.2.0-cdh5.8.5

Project	groupId	artifactId	version
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.8.5
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.8.5
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.8.5
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.8.5
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.8.5
	org.apache.pig	piggybank	0.12.0-cdh5.8.5
	org.apache.pig	pigsmoke	0.12.0-cdh5.8.5
	org.apache.pig	pigunit	0.12.0-cdh5.8.5
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.8.5
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.8.5
	org.apache.flume	flume-ng-core	1.6.0-cdh5.8.5
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.8.5
	org.apache.flume	flume-ng-node	1.6.0-cdh5.8.5
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.8.5
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.8.5
	org.apache.flume	flume-tools	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.8.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sources	flume-jms-source	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.8.5
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.8.5

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.8.5
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.6.0-cdh5.8.5
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-core	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-examples	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.8.5.oozie-4.1.0-cdh5.8.5
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.8.5.oozie-4.1.0-cdh5.8.5
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.8.5.oozie-4.1.0-cdh5.8.5
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.8.5.oozie-4.1.0-cdh5.8.5
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.8.5.oozie-4.1.0-cdh5.8.5
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.8.5.oozie-4.1.0-cdh5.8.5
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.8.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.8.5
	org.apache.oozie	oozie-tools	4.1.0-cdh5.8.5
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.8.5
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.8.5
	org.apache.mahout	mahout-core	0.9-cdh5.8.5
	org.apache.mahout	mahout-examples	0.9-cdh5.8.5
	org.apache.mahout	mahout-integration	0.9-cdh5.8.5
	org.apache.mahout	mahout-math	0.9-cdh5.8.5
	org.apache.mahout	mahout-math-scala	0.9-cdh5.8.5
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-chef	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-cli	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-core	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-druid	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-examples	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.8.5

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-hama	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-pig	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-solr	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.8.5
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.8.5
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.8.5
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.8.5
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.8.5
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.8.5
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.8.5
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.8.5
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.8.5
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.8.5
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.8.5
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.8.5
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.8.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	test	1.99.5-cdh5.8.5
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.8.5
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.8.5
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.8.5
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.8.5
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.8.5
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.8.5
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.8.5
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.8.5
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.8.5
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.8.5

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-dist	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.8.5
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.8.5
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.8.5
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.8.5
	com.twitter	parquet-cascading	1.5.0-cdh5.8.5
	com.twitter	parquet-column	1.5.0-cdh5.8.5

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-common	1.5.0-cdh5.8.5
	com.twitter	parquet-encoding	1.5.0-cdh5.8.5
	com.twitter	parquet-format	2.1.0-cdh5.8.5
	com.twitter	parquet-generator	1.5.0-cdh5.8.5
	com.twitter	parquet-hadoop	1.5.0-cdh5.8.5
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.8.5
	com.twitter	parquet-jackson	1.5.0-cdh5.8.5
	com.twitter	parquet-pig	1.5.0-cdh5.8.5
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.8.5
	com.twitter	parquet-protobuf	1.5.0-cdh5.8.5
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.8.5
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.8.5
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.8.5
	com.twitter	parquet-thrift	1.5.0-cdh5.8.5
	com.twitter	parquet-tools	1.5.0-cdh5.8.5
Llama	com.cloudera.llama	llama	1.0.0-cdh5.8.5
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.8.5
	com.databricks	spark-avro_2.10	1.1.0-cdh5.8.5
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.8.5

Project	groupId	artifactId	version
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.8.5
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.8.5
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.8.5
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.8.5
	org.apache.crunch	crunch-core	0.11.0-cdh5.8.5
	org.apache.crunch	crunch-examples	0.11.0-cdh5.8.5
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.8.5
	org.apache.crunch	crunch-hive	0.11.0-cdh5.8.5
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.8.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-spark	0.11.0-cdh5.8.5
	org.apache.crunch	crunch-test	0.11.0-cdh5.8.5
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.8.5
	org.apache.avro	avro-compiler	1.7.6-cdh5.8.5
	org.apache.avro	avro-ipc	1.7.6-cdh5.8.5
	org.apache.avro	avro-mapred	1.7.6-cdh5.8.5
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.8.5
	org.apache.avro	avro-protobuf	1.7.6-cdh5.8.5
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.8.5
	org.apache.avro	avro-thrift	1.7.6-cdh5.8.5
	org.apache.avro	avro-tools	1.7.6-cdh5.8.5
	org.apache.avro	trevni-avro	1.7.6-cdh5.8.5
	org.apache.avro	trevni-core	1.7.6-cdh5.8.5
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.8.5
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.8.5
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.8.5
	org.kitesdk	kite-data-hive	1.0.0-cdh5.8.5
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.8.5
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.8.5
	org.kitesdk	kite-data-s3	1.0.0-cdh5.8.5
	org.kitesdk	kite-data-spark	1.0.0-cdh5.8.5
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.8.5
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.8.5
	org.kitesdk	kite-minicluster	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.8.5

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.8.5
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.8.5
	org.kitesdk	kite-tools	1.0.0-cdh5.8.5
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-classification	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-core	4.10.3-cdh5.8.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-demo	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-facet	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-join	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-memory	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-misc	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-queries	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.8.5
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.8.5
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.8.5
	org.apache.solr	solr-cell	4.10.3-cdh5.8.5
	org.apache.solr	solr-clustering	4.10.3-cdh5.8.5
	org.apache.solr	solr-core	4.10.3-cdh5.8.5
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.8.5
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.8.5

Project	groupId	artifactId	version
	org.apache.solr	solr-langid	4.10.3-cdh5.8.5
	org.apache.solr	solr-security-util	4.10.3-cdh5.8.5
	org.apache.solr	solr-solrj	4.10.3-cdh5.8.5
	org.apache.solr	solr-test-framework	4.10.3-cdh5.8.5
	org.apache.solr	solr-uima	4.10.3-cdh5.8.5
	org.apache.solr	solr-velocity	4.10.3-cdh5.8.5
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.8.5
	com.cloudera.search	search-mr	1.0.0-cdh5.8.5
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.8.5
	com.ngdata	hbase-indexer-cli	1.5-cdh5.8.5
	com.ngdata	hbase-indexer-common	1.5-cdh5.8.5
	com.ngdata	hbase-indexer-demo	1.5-cdh5.8.5
	com.ngdata	hbase-indexer-dist	1.5-cdh5.8.5
	com.ngdata	hbase-indexer-engine	1.5-cdh5.8.5
	com.ngdata	hbase-indexer-model	1.5-cdh5.8.5
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.8.5
	com.ngdata	hbase-indexer-mr	1.5-cdh5.8.5
	com.ngdata	hbase-indexer-server	1.5-cdh5.8.5
	com.ngdata	hbase-sep-api	1.5-cdh5.8.5
	com.ngdata	hbase-sep-demo	1.5-cdh5.8.5
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.8.5
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.8.5
	com.ngdata	hbase-sep-tools	1.5-cdh5.8.5

CDH 5.8.4

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.8.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.8.4

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.8.4
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.8.4
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.8.4
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.8.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.8.4
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.8.4
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.8.4
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.8.4
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.8.4
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.8.4
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.8.4
	org.apache.hive	hive-ant	1.1.0-cdh5.8.4
	org.apache.hive	hive-beeline	1.1.0-cdh5.8.4
	org.apache.hive	hive-cli	1.1.0-cdh5.8.4
	org.apache.hive	hive-common	1.1.0-cdh5.8.4
	org.apache.hive	hive-contrib	1.1.0-cdh5.8.4
	org.apache.hive	hive-exec	1.1.0-cdh5.8.4
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.8.4
	org.apache.hive	hive-hwi	1.1.0-cdh5.8.4
	org.apache.hive	hive-jdbc	1.1.0-cdh5.8.4
	org.apache.hive	hive-metastore	1.1.0-cdh5.8.4
	org.apache.hive	hive-serde	1.1.0-cdh5.8.4
	org.apache.hive	hive-service	1.1.0-cdh5.8.4

Project	groupId	artifactId	version
	org.apache.hive	hive-shims	1.1.0-cdh5.8.4
	org.apache.hive	hive-testutils	1.1.0-cdh5.8.4
	org.apache.hive	spark-client	1.1.0-cdh5.8.4
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.8.4
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.8.4
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.8.4
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.8.4
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.8.4
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.8.4
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.8.4
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.8.4
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.8.4
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-client	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-common	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-examples	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-it	1.2.0-cdh5.8.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-rest	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-server	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-shell	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-spark	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.8.4
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.8.4
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.8.4
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.8.4
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.8.4
	org.apache.pig	piggybank	0.12.0-cdh5.8.4
	org.apache.pig	pigsmoke	0.12.0-cdh5.8.4
	org.apache.pig	pigunit	0.12.0-cdh5.8.4
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.8.4
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.8.4
	org.apache.flume	flume-ng-core	1.6.0-cdh5.8.4

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.8.4
	org.apache.flume	flume-ng-node	1.6.0-cdh5.8.4
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.8.4
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.8.4
	org.apache.flume	flume-tools	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.8.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sources	flume-jms-source	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.8.4
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.8.4
	org.apache.flume.flume-shared	flume-shared-kafka-test	1.6.0-cdh5.8.4
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-core	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-examples	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.8.4.oozie-4.1.0-cdh5.8.4
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.8.4.oozie-4.1.0-cdh5.8.4
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.8.4.oozie-4.1.0-cdh5.8.4

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.8.4.oozie-4.1.0-cdh5.8.4
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.8.4.oozie-4.1.0-cdh5.8.4
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.8.4.oozie-4.1.0-cdh5.8.4
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.8.4
	org.apache.oozie	oozie-tools	4.1.0-cdh5.8.4
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.8.4
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.8.4
	org.apache.mahout	mahout-core	0.9-cdh5.8.4
	org.apache.mahout	mahout-examples	0.9-cdh5.8.4
	org.apache.mahout	mahout-integration	0.9-cdh5.8.4
	org.apache.mahout	mahout-math	0.9-cdh5.8.4
	org.apache.mahout	mahout-math-scala	0.9-cdh5.8.4

Version and Download Information

Project	groupId	artifactId	version
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-chef	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-cli	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-core	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-druid	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-examples	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-hama	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-pig	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-solr	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.8.4
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.8.4
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.8.4

Project	groupId	artifactId	version
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.8.4
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.8.4
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.8.4
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.8.4
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.8.4
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.8.4
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.8.4
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.8.4
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.8.4
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.8.4
	org.apache.sqoop	test	1.99.5-cdh5.8.4
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.8.4
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.8.4
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.8.4
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.8.4
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.8.4
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.8.4
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.8.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.8.4
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.8.4
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-dist	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.8.4

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.8.4
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.8.4
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.8.4
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.8.4
	com.twitter	parquet-cascading	1.5.0-cdh5.8.4
	com.twitter	parquet-column	1.5.0-cdh5.8.4
	com.twitter	parquet-common	1.5.0-cdh5.8.4
	com.twitter	parquet-encoding	1.5.0-cdh5.8.4
	com.twitter	parquet-format	2.1.0-cdh5.8.4
	com.twitter	parquet-generator	1.5.0-cdh5.8.4
	com.twitter	parquet-hadoop	1.5.0-cdh5.8.4
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.8.4
	com.twitter	parquet-jackson	1.5.0-cdh5.8.4
	com.twitter	parquet-pig	1.5.0-cdh5.8.4
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.8.4
	com.twitter	parquet-protobuf	1.5.0-cdh5.8.4
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.8.4
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.8.4
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.8.4
	com.twitter	parquet-thrift	1.5.0-cdh5.8.4
	com.twitter	parquet-tools	1.5.0-cdh5.8.4
Llama	com.cloudera.llama	llama	1.0.0-cdh5.8.4

Version and Download Information

Project	groupId	artifactId	version
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.8.4

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.8.4
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.8.4
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.8.4
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.8.4
	org.apache.crunch	crunch-core	0.11.0-cdh5.8.4
	org.apache.crunch	crunch-examples	0.11.0-cdh5.8.4
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.8.4
	org.apache.crunch	crunch-hive	0.11.0-cdh5.8.4
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.8.4
	org.apache.crunch	crunch-spark	0.11.0-cdh5.8.4
	org.apache.crunch	crunch-test	0.11.0-cdh5.8.4
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.8.4
	org.apache.avro	avro-compiler	1.7.6-cdh5.8.4
	org.apache.avro	avro-ipc	1.7.6-cdh5.8.4
	org.apache.avro	avro-mapred	1.7.6-cdh5.8.4
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.8.4
	org.apache.avro	avro-protobuf	1.7.6-cdh5.8.4
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.8.4
	org.apache.avro	avro-thrift	1.7.6-cdh5.8.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	avro-tools	1.7.6-cdh5.8.4
	org.apache.avro	trevni-avro	1.7.6-cdh5.8.4
	org.apache.avro	trevni-core	1.7.6-cdh5.8.4
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.8.4
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.8.4
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.8.4
	org.kitesdk	kite-data-hive	1.0.0-cdh5.8.4
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.8.4
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.8.4
	org.kitesdk	kite-data-s3	1.0.0-cdh5.8.4
	org.kitesdk	kite-data-spark	1.0.0-cdh5.8.4
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.8.4
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.8.4
	org.kitesdk	kite-minicluster	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.8.4
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.8.4
	org.kitesdk	kite-tools	1.0.0-cdh5.8.4

Project	groupId	artifactId	version
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-classification	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-core	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-demo	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-facet	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-join	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-memory	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-misc	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-queries	4.10.3-cdh5.8.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.8.4
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.8.4
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.8.4
	org.apache.solr	solr-cell	4.10.3-cdh5.8.4
	org.apache.solr	solr-clustering	4.10.3-cdh5.8.4
	org.apache.solr	solr-core	4.10.3-cdh5.8.4
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.8.4
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.8.4
	org.apache.solr	solr-langid	4.10.3-cdh5.8.4
	org.apache.solr	solr-security-util	4.10.3-cdh5.8.4
	org.apache.solr	solr-solrj	4.10.3-cdh5.8.4
	org.apache.solr	solr-test-framework	4.10.3-cdh5.8.4
	org.apache.solr	solr-uima	4.10.3-cdh5.8.4
	org.apache.solr	solr-velocity	4.10.3-cdh5.8.4
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.8.4
	com.cloudera.search	search-mr	1.0.0-cdh5.8.4
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.8.4
	com.ngdata	hbase-indexer-cli	1.5-cdh5.8.4

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-common	1.5-cdh5.8.4
	com.ngdata	hbase-indexer-demo	1.5-cdh5.8.4
	com.ngdata	hbase-indexer-dist	1.5-cdh5.8.4
	com.ngdata	hbase-indexer-engine	1.5-cdh5.8.4
	com.ngdata	hbase-indexer-model	1.5-cdh5.8.4
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.8.4
	com.ngdata	hbase-indexer-mr	1.5-cdh5.8.4
	com.ngdata	hbase-indexer-server	1.5-cdh5.8.4
	com.ngdata	hbase-sep-api	1.5-cdh5.8.4
	com.ngdata	hbase-sep-demo	1.5-cdh5.8.4
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.8.4
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.8.4
	com.ngdata	hbase-sep-tools	1.5-cdh5.8.4

CDH 5.8.3

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.8.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.8.3

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.8.3
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.8.3
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.8.3
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.8.3
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.8.3
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.8.3
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.8.3
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.8.3
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.8.3
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.8.3

Version and Download Information

Project	groupId	artifactId	version
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.8.3
	org.apache.hive	hive-ant	1.1.0-cdh5.8.3
	org.apache.hive	hive-beeline	1.1.0-cdh5.8.3
	org.apache.hive	hive-cli	1.1.0-cdh5.8.3
	org.apache.hive	hive-common	1.1.0-cdh5.8.3
	org.apache.hive	hive-contrib	1.1.0-cdh5.8.3
	org.apache.hive	hive-exec	1.1.0-cdh5.8.3
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.8.3
	org.apache.hive	hive-hwi	1.1.0-cdh5.8.3
	org.apache.hive	hive-jdbc	1.1.0-cdh5.8.3
	org.apache.hive	hive-metastore	1.1.0-cdh5.8.3
	org.apache.hive	hive-serde	1.1.0-cdh5.8.3
	org.apache.hive	hive-service	1.1.0-cdh5.8.3
	org.apache.hive	hive-shims	1.1.0-cdh5.8.3
	org.apache.hive	hive-testutils	1.1.0-cdh5.8.3
	org.apache.hive	spark-client	1.1.0-cdh5.8.3
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.8.3
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.8.3
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.8.3
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.8.3
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.8.3

Project	groupId	artifactId	version
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.8.3
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.8.3
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.8.3
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.8.3
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-client	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-common	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-examples	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-it	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-rest	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-server	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.8.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shell	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-spark	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.8.3
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.8.3
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.8.3
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.8.3
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.8.3
	org.apache.pig	piggybank	0.12.0-cdh5.8.3
	org.apache.pig	pigsmoke	0.12.0-cdh5.8.3
	org.apache.pig	pigunit	0.12.0-cdh5.8.3
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.8.3
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.8.3
	org.apache.flume	flume-ng-core	1.6.0-cdh5.8.3
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.8.3
	org.apache.flume	flume-ng-node	1.6.0-cdh5.8.3
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.8.3
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.8.3
	org.apache.flume	flume-tools	1.6.0-cdh5.8.3
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.8.3
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.8.3

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.8.3
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.8.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.8.3
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.8.3
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.8.3
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-core	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-examples	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.8.3.oozie-4.1.0-cdh5.8.3
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.8.3.oozie-4.1.0-cdh5.8.3
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.8.3.oozie-4.1.0-cdh5.8.3
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.8.3.oozie-4.1.0-cdh5.8.3
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.8.3.oozie-4.1.0-cdh5.8.3
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.8.3.oozie-4.1.0-cdh5.8.3
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.8.3

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.8.3
	org.apache.oozie	oozie-tools	4.1.0-cdh5.8.3
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.8.3
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.8.3
	org.apache.mahout	mahout-core	0.9-cdh5.8.3
	org.apache.mahout	mahout-examples	0.9-cdh5.8.3
	org.apache.mahout	mahout-integration	0.9-cdh5.8.3
	org.apache.mahout	mahout-math	0.9-cdh5.8.3
	org.apache.mahout	mahout-math-scala	0.9-cdh5.8.3
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-chef	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-cli	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-core	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-druid	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.8.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-examples	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-hama	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-pig	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-solr	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.8.3
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.8.3
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.8.3
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.8.3
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.8.3
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.8.3
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.8.3
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.8.3
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.8.3
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.8.3
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.8.3

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.8.3
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.8.3
	org.apache.sqoop	test	1.99.5-cdh5.8.3
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.8.3
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.8.3
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.8.3
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.8.3
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.8.3
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.8.3
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.8.3
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.8.3
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.8.3
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.8.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-dist	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.8.3
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.8.3
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.8.3

Project	groupId	artifactId	version
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.8.3
	com.twitter	parquet-cascading	1.5.0-cdh5.8.3
	com.twitter	parquet-column	1.5.0-cdh5.8.3
	com.twitter	parquet-common	1.5.0-cdh5.8.3
	com.twitter	parquet-encoding	1.5.0-cdh5.8.3
	com.twitter	parquet-format	2.1.0-cdh5.8.3
	com.twitter	parquet-generator	1.5.0-cdh5.8.3
	com.twitter	parquet-hadoop	1.5.0-cdh5.8.3
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.8.3
	com.twitter	parquet-jackson	1.5.0-cdh5.8.3
	com.twitter	parquet-pig	1.5.0-cdh5.8.3
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.8.3
	com.twitter	parquet-protobuf	1.5.0-cdh5.8.3
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.8.3
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.8.3
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.8.3
	com.twitter	parquet-thrift	1.5.0-cdh5.8.3
	com.twitter	parquet-tools	1.5.0-cdh5.8.3
Llama	com.cloudera.llama	llama	1.0.0-cdh5.8.3
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.8.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.8.3
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.8.3
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.8.3
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.8.3
	org.apache.crunch	crunch-core	0.11.0-cdh5.8.3
	org.apache.crunch	crunch-examples	0.11.0-cdh5.8.3
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.8.3

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hive	0.11.0-cdh5.8.3
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.8.3
	org.apache.crunch	crunch-spark	0.11.0-cdh5.8.3
	org.apache.crunch	crunch-test	0.11.0-cdh5.8.3
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.8.3
	org.apache.avro	avro-compiler	1.7.6-cdh5.8.3
	org.apache.avro	avro-ipc	1.7.6-cdh5.8.3
	org.apache.avro	avro-mapred	1.7.6-cdh5.8.3
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.8.3
	org.apache.avro	avro-protobuf	1.7.6-cdh5.8.3
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.8.3
	org.apache.avro	avro-thrift	1.7.6-cdh5.8.3
	org.apache.avro	avro-tools	1.7.6-cdh5.8.3
	org.apache.avro	trevni-avro	1.7.6-cdh5.8.3
	org.apache.avro	trevni-core	1.7.6-cdh5.8.3
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.8.3
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.8.3
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.8.3
	org.kitesdk	kite-data-hive	1.0.0-cdh5.8.3
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.8.3
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.8.3
	org.kitesdk	kite-data-s3	1.0.0-cdh5.8.3
	org.kitesdk	kite-data-spark	1.0.0-cdh5.8.3
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.8.3
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.8.3
	org.kitesdk	kite-minicluster	1.0.0-cdh5.8.3

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.8.3
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.8.3
	org.kitesdk	kite-tools	1.0.0-cdh5.8.3
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-classification	4.10.3-cdh5.8.3

Project	groupId	artifactId	version
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-core	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-demo	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-facet	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-join	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-memory	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-misc	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-queries	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.8.3
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.8.3
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.8.3
	org.apache.solr	solr-cell	4.10.3-cdh5.8.3
	org.apache.solr	solr-clustering	4.10.3-cdh5.8.3
	org.apache.solr	solr-core	4.10.3-cdh5.8.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.8.3
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.8.3
	org.apache.solr	solr-langid	4.10.3-cdh5.8.3
	org.apache.solr	solr-security-util	4.10.3-cdh5.8.3
	org.apache.solr	solr-solrj	4.10.3-cdh5.8.3
	org.apache.solr	solr-test-framework	4.10.3-cdh5.8.3
	org.apache.solr	solr-uima	4.10.3-cdh5.8.3
	org.apache.solr	solr-velocity	4.10.3-cdh5.8.3
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.8.3
	com.cloudera.search	search-mr	1.0.0-cdh5.8.3
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.8.3
	com.ngdata	hbase-indexer-cli	1.5-cdh5.8.3
	com.ngdata	hbase-indexer-common	1.5-cdh5.8.3
	com.ngdata	hbase-indexer-demo	1.5-cdh5.8.3
	com.ngdata	hbase-indexer-dist	1.5-cdh5.8.3
	com.ngdata	hbase-indexer-engine	1.5-cdh5.8.3
	com.ngdata	hbase-indexer-model	1.5-cdh5.8.3
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.8.3
	com.ngdata	hbase-indexer-mr	1.5-cdh5.8.3
	com.ngdata	hbase-indexer-server	1.5-cdh5.8.3
	com.ngdata	hbase-sep-api	1.5-cdh5.8.3
	com.ngdata	hbase-sep-demo	1.5-cdh5.8.3
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.8.3
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.8.3
	com.ngdata	hbase-sep-tools	1.5-cdh5.8.3

CDH 5.8.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.8.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.8.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.8.2

Project	groupId	artifactId	version
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.8.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.8.2
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.8.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.8.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.8.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.8.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.8.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.8.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.8.2
	org.apache.hive	hive-ant	1.1.0-cdh5.8.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.8.2
	org.apache.hive	hive-cli	1.1.0-cdh5.8.2
	org.apache.hive	hive-common	1.1.0-cdh5.8.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.8.2
	org.apache.hive	hive-exec	1.1.0-cdh5.8.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.8.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.8.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.8.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-metastore	1.1.0-cdh5.8.2
	org.apache.hive	hive-serde	1.1.0-cdh5.8.2
	org.apache.hive	hive-service	1.1.0-cdh5.8.2
	org.apache.hive	hive-shims	1.1.0-cdh5.8.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.8.2
	org.apache.hive	spark-client	1.1.0-cdh5.8.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.8.2
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.8.2
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.8.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.8.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.8.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.8.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.8.2
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.8.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.8.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.8.2

Project	groupId	artifactId	version
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-it	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-rest	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-server	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.8.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.8.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.8.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.8.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.8.2
	org.apache.pig	piggybank	0.12.0-cdh5.8.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.8.2
	org.apache.pig	pigunit	0.12.0-cdh5.8.2

Version and Download Information

Project	groupId	artifactId	version
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.8.2
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.8.2
	org.apache.flume	flume-ng-core	1.6.0-cdh5.8.2
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.8.2
	org.apache.flume	flume-ng-node	1.6.0-cdh5.8.2
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.8.2
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.8.2
	org.apache.flume	flume-tools	1.6.0-cdh5.8.2
	org.apache.flume. flume-ng-channels	flume-file-channel	1.6.0-cdh5.8.2
	org.apache.flume. flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.8.2
	org.apache.flume. flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.8.2
	org.apache.flume. flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.8.2
	org.apache.flume. flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.8.2
	org.apache.flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.8.2
	org.apache.flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.8.2
	org.apache.flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.8.2
	org.apache.flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.8.2

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.8.2
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.8.2
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.8.2
	org.apache. oozie	oozie-core	4.1.0-cdh5.8.2
	org.apache. oozie	oozie-examples	4.1.0-cdh5.8.2
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.8. 2.oozie-4.1. 0-cdh5.8.2
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.8.

Version and Download Information

Project	groupId	artifactId	version
			2.oozie-4.1.0-cdh5.8.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.8.2.oozie-4.1.0-cdh5.8.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.8.2.oozie-4.1.0-cdh5.8.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.8.2.oozie-4.1.0-cdh5.8.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.8.2.oozie-4.1.0-cdh5.8.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.8.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.8.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.8.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.8.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.8.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.8.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.8.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.8.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.8.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.8.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.8.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.8.2
	org.apache.mahout	mahout-core	0.9-cdh5.8.2
	org.apache.mahout	mahout-examples	0.9-cdh5.8.2

Project	groupId	artifactId	version
	org.apache.mahout	mahout-integration	0.9-cdh5.8.2
	org.apache.mahout	mahout-math	0.9-cdh5.8.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.8.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.8.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.8.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.8.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.8.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.8.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.8.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.8.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.8.2
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.8.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.8.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.8.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.8.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.8.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.8.2
	org.apache.sqoop	test	1.99.5-cdh5.8.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.8.2
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.8.2
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.8.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.8.2
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.8.2

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.8.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.8.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.8.2
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.8.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.8.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.8.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.8.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.8.2
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.8.2
	com.twitter	parquet-cascading	1.5.0-cdh5.8.2
	com.twitter	parquet-column	1.5.0-cdh5.8.2
	com.twitter	parquet-common	1.5.0-cdh5.8.2
	com.twitter	parquet-encoding	1.5.0-cdh5.8.2
	com.twitter	parquet-format	2.1.0-cdh5.8.2
	com.twitter	parquet-generator	1.5.0-cdh5.8.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.8.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.8.2
	com.twitter	parquet-jackson	1.5.0-cdh5.8.2
	com.twitter	parquet-pig	1.5.0-cdh5.8.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.8.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.8.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.8.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.8.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.8.2
	com.twitter	parquet-thrift	1.5.0-cdh5.8.2

Project	groupId	artifactId	version
	com.twitter	parquet-tools	1.5.0-cdh5.8.2
Llama	com.cloudera. llama	llama	1.0.0-cdh5.8.2
Apache Spark	org.apache. spark	spark-assembly_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-bagel_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-catalyst_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-core_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-graphx_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-hive_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-launcher_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-mllib_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-network-common_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-network-shuffle_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-network-yarn_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-repl_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-sql_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-streaming-flume_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-streaming-kafka_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.8.2
	org.apache. spark	spark-streaming-twitter_2.10	1.6.0-cdh5.8.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.8.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.8.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.8.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.8.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.8.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.8.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.8.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.8.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.8.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.8.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.8.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.8.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.8.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.8.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.8.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.8.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.8.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.8.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.8.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.8.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.8.2

Project	groupId	artifactId	version
	org.apache.avro	avro-thrift	1.7.6-cdh5.8.2
	org.apache.avro	avro-tools	1.7.6-cdh5.8.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.8.2
	org.apache.avro	trevni-core	1.7.6-cdh5.8.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.8.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.8.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.8.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.8.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.8.2
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.8.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.8.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.8.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.8.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.8.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.8.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.8.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-tools	1.0.0-cdh5.8.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.8.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.8.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.8.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.8.2
	org.apache.solr	solr-cell	4.10.3-cdh5.8.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.8.2
	org.apache.solr	solr-core	4.10.3-cdh5.8.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.8.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.8.2
	org.apache.solr	solr-langid	4.10.3-cdh5.8.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.8.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.8.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.8.2
	org.apache.solr	solr-uima	4.10.3-cdh5.8.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.8.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.8.2
	com.cloudera.search	search-mr	1.0.0-cdh5.8.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.8.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.8.2

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-common	1.5-cdh5.8.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.8.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.8.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.8.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.8.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.8.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.8.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.8.2
	com.ngdata	hbase-sep-api	1.5-cdh5.8.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.8.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.8.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.8.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.8.2

CDH 5.8.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.8.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.8.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.8.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.8.0
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.8.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.8.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.8.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.8.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.8.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.8.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.8.0
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.8.0

Project	groupId	artifactId	version
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.8.0
	org.apache.hive	hive-ant	1.1.0-cdh5.8.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.8.0
	org.apache.hive	hive-cli	1.1.0-cdh5.8.0
	org.apache.hive	hive-common	1.1.0-cdh5.8.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.8.0
	org.apache.hive	hive-exec	1.1.0-cdh5.8.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.8.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.8.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.8.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.8.0
	org.apache.hive	hive-serde	1.1.0-cdh5.8.0
	org.apache.hive	hive-service	1.1.0-cdh5.8.0
	org.apache.hive	hive-shims	1.1.0-cdh5.8.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.8.0
	org.apache.hive	spark-client	1.1.0-cdh5.8.0
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.8.0
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.8.0
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.8.0
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.8.0
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.8.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.8.0
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.8.0
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.8.0
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.8.0
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-client	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-common	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-examples	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-it	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-rest	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-server	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.8.0

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shell	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-spark	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.8.0
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.8.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.8.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.8.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.8.0
	org.apache.pig	piggybank	0.12.0-cdh5.8.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.8.0
	org.apache.pig	pigunit	0.12.0-cdh5.8.0
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.8.0
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.8.0
	org.apache.flume	flume-ng-core	1.6.0-cdh5.8.0
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.8.0
	org.apache.flume	flume-ng-node	1.6.0-cdh5.8.0
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.8.0
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.8.0
	org.apache.flume	flume-tools	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.8.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sources	flume-jms-source	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.6.0-cdh5.8.0

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.8.0
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.8.0
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-core	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-examples	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.8.0.oozie-4.1.0-cdh5.8.0
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.8.0.oozie-4.1.0-cdh5.8.0
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.8.0.oozie-4.1.0-cdh5.8.0
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.8.0.oozie-4.1.0-cdh5.8.0
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.8.0.oozie-4.1.0-cdh5.8.0
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.8.0.oozie-4.1.0-cdh5.8.0
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.8.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.8.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.8.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.8.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.8.0
	org.apache.mahout	mahout-core	0.9-cdh5.8.0
	org.apache.mahout	mahout-examples	0.9-cdh5.8.0
	org.apache.mahout	mahout-integration	0.9-cdh5.8.0
	org.apache.mahout	mahout-math	0.9-cdh5.8.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.8.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.8.0

Project	groupId	artifactId	version
	org.apache.whirr	whirr-examples	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.8.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.8.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.8.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.8.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.8.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.8.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.8.0
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.8.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.8.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.8.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.8.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.8.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.8.0
	org.apache.sqoop	test	1.99.5-cdh5.8.0
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.8.0
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.8.0
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.8.0
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.8.0
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.8.0
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.8.0
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.8.0
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.8.0
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.8.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-binding-kafka	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.8.0

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-core-model-kafka	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-policy-kafka	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-tests-kafka	1.5.1-cdh5.8.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.8.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.8.0

Version and Download Information

Project	groupId	artifactId	version
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.8.0
	com.twitter	parquet-cascading	1.5.0-cdh5.8.0
	com.twitter	parquet-column	1.5.0-cdh5.8.0
	com.twitter	parquet-common	1.5.0-cdh5.8.0
	com.twitter	parquet-encoding	1.5.0-cdh5.8.0
	com.twitter	parquet-format	2.1.0-cdh5.8.0
	com.twitter	parquet-generator	1.5.0-cdh5.8.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.8.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.8.0
	com.twitter	parquet-jackson	1.5.0-cdh5.8.0
	com.twitter	parquet-pig	1.5.0-cdh5.8.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.8.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.8.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.8.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.8.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.8.0
	com.twitter	parquet-thrift	1.5.0-cdh5.8.0
	com.twitter	parquet-tools	1.5.0-cdh5.8.0
Llama	com.cloudera.llama	llama	1.0.0-cdh5.8.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.8.0

Project	groupId	artifactId	version
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.8.0
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.8.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.8.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.8.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.8.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.8.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.8.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hive	0.11.0-cdh5.8.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.8.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.8.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.8.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.8.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.8.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.8.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.8.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.8.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.8.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.8.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.8.0
	org.apache.avro	avro-tools	1.7.6-cdh5.8.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.8.0
	org.apache.avro	trevni-core	1.7.6-cdh5.8.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.8.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.8.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.8.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.8.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.8.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.8.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.8.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.8.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.8.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.8.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.8.0

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.8.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.8.0
	org.kitesdk	kite-tools	1.0.0-cdh5.8.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.8.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.8.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.8.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.8.0
	org.apache.solr	solr-cell	4.10.3-cdh5.8.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.8.0
	org.apache.solr	solr-core	4.10.3-cdh5.8.0

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.8.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.8.0
	org.apache.solr	solr-langid	4.10.3-cdh5.8.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.8.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.8.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.8.0
	org.apache.solr	solr-uima	4.10.3-cdh5.8.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.8.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.8.0
	com.cloudera.search	search-mr	1.0.0-cdh5.8.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.8.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.8.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.8.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.8.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.8.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.8.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.8.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.8.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.8.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.8.0
	com.ngdata	hbase-sep-api	1.5-cdh5.8.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.8.0
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.8.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.8.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.8.0

Maven Artifacts for CDH 5.7.x Releases CDH 5.7.6

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Version and Download Information

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.7.6

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.7.6
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.7.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.7.6
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.7.6
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.7.6
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.7.6
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.7.6
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.7.6
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.7.6
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.7.6
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.7.6
	org.apache.hive	hive-ant	1.1.0-cdh5.7.6
	org.apache.hive	hive-beeline	1.1.0-cdh5.7.6
	org.apache.hive	hive-cli	1.1.0-cdh5.7.6
	org.apache.hive	hive-common	1.1.0-cdh5.7.6
	org.apache.hive	hive-contrib	1.1.0-cdh5.7.6
	org.apache.hive	hive-exec	1.1.0-cdh5.7.6
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.7.6
	org.apache.hive	hive-hwi	1.1.0-cdh5.7.6
	org.apache.hive	hive-jdbc	1.1.0-cdh5.7.6

Project	groupId	artifactId	version
	org.apache.hive	hive-metastore	1.1.0-cdh5.7.6
	org.apache.hive	hive-serde	1.1.0-cdh5.7.6
	org.apache.hive	hive-service	1.1.0-cdh5.7.6
	org.apache.hive	hive-shims	1.1.0-cdh5.7.6
	org.apache.hive	hive-testutils	1.1.0-cdh5.7.6
	org.apache.hive	spark-client	1.1.0-cdh5.7.6
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.7.6
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.7.6
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.7.6
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.7.6
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.7.6
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.7.6
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.7.6
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.7.6
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.7.6
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-client	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-common	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-examples	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.7.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-it	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-rest	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-server	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-shell	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-spark	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.7.6
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.7.6
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.7.6
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.7.6
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.7.6
	org.apache.pig	piggybank	0.12.0-cdh5.7.6
	org.apache.pig	pigsmoke	0.12.0-cdh5.7.6
	org.apache.pig	pigunit	0.12.0-cdh5.7.6

Project	groupId	artifactId	version
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.7.6
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.7.6
	org.apache.flume	flume-ng-core	1.6.0-cdh5.7.6
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.7.6
	org.apache.flume	flume-ng-node	1.6.0-cdh5.7.6
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.7.6
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.7.6
	org.apache.flume	flume-tools	1.6.0-cdh5.7.6
	org.apache.flume. flume-ng-channels	flume-file-channel	1.6.0-cdh5.7.6
	org.apache.flume. flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.7.6
	org.apache.flume. flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.7.6
	org.apache.flume. flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.7.6
	org.apache.flume. flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.7.6
	org.apache.flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.7.6
	org.apache.flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.7.6
	org.apache.flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.7.6
	org.apache.flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.7.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sources	flume-jms-source	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.7.6
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.7.6
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-core	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-examples	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.7.6.oozie-4.1.0-cdh5.7.6
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.7.

Project	groupId	artifactId	version
			6.oozie-4.1.0-cdh5.7.6
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.7.6.oozie-4.1.0-cdh5.7.6
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.7.6.oozie-4.1.0-cdh5.7.6
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.7.6.oozie-4.1.0-cdh5.7.6
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.7.6.oozie-4.1.0-cdh5.7.6
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.7.6
	org.apache.oozie	oozie-tools	4.1.0-cdh5.7.6
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.7.6
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.7.6
	org.apache.mahout	mahout-core	0.9-cdh5.7.6
	org.apache.mahout	mahout-examples	0.9-cdh5.7.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.mahout	mahout-integration	0.9-cdh5.7.6
	org.apache.mahout	mahout-math	0.9-cdh5.7.6
	org.apache.mahout	mahout-math-scala	0.9-cdh5.7.6
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-chef	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-cli	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-core	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-druid	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-examples	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-hama	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-pig	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-solr	0.9.0-cdh5.7.6

Project	groupId	artifactId	version
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.7.6
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.7.6
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.7.6
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.7.6
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.7.6
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.7.6
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.7.6
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.7.6
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.7.6
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.7.6
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.7.6
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.7.6
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.7.6
	org.apache.sqoop	test	1.99.5-cdh5.7.6
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.7.6
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.7.6
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.7.6
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.7.6
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.7.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.7.6
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.7.6
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.7.6
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.7.6
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-dist	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.7.6

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.7.6
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.7.6
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.7.6
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.7.6
	com.twitter	parquet-cascading	1.5.0-cdh5.7.6
	com.twitter	parquet-column	1.5.0-cdh5.7.6
	com.twitter	parquet-common	1.5.0-cdh5.7.6
	com.twitter	parquet-encoding	1.5.0-cdh5.7.6
	com.twitter	parquet-format	2.1.0-cdh5.7.6
	com.twitter	parquet-generator	1.5.0-cdh5.7.6
	com.twitter	parquet-hadoop	1.5.0-cdh5.7.6
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.7.6
	com.twitter	parquet-jackson	1.5.0-cdh5.7.6
	com.twitter	parquet-pig	1.5.0-cdh5.7.6
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.7.6
	com.twitter	parquet-protobuf	1.5.0-cdh5.7.6
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.7.6
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.7.6
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.7.6
	com.twitter	parquet-thrift	1.5.0-cdh5.7.6
	com.twitter	parquet-tools	1.5.0-cdh5.7.6
Llama	com.cloudera.llama	llama	1.0.0-cdh5.7.6
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.7.6
	com.databricks	spark-avro_2.10	1.1.0-cdh5.7.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.7.6

Project	groupId	artifactId	version
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.7.6
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.7.6
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.7.6
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.7.6
	org.apache.crunch	crunch-core	0.11.0-cdh5.7.6
	org.apache.crunch	crunch-examples	0.11.0-cdh5.7.6
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.7.6
	org.apache.crunch	crunch-hive	0.11.0-cdh5.7.6
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.7.6
	org.apache.crunch	crunch-spark	0.11.0-cdh5.7.6
	org.apache.crunch	crunch-test	0.11.0-cdh5.7.6
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.7.6
	org.apache.avro	avro-compiler	1.7.6-cdh5.7.6
	org.apache.avro	avro-ipc	1.7.6-cdh5.7.6
	org.apache.avro	avro-mapred	1.7.6-cdh5.7.6
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.7.6
	org.apache.avro	avro-protobuf	1.7.6-cdh5.7.6
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.7.6
	org.apache.avro	avro-thrift	1.7.6-cdh5.7.6
	org.apache.avro	avro-tools	1.7.6-cdh5.7.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	trevni-avro	1.7.6-cdh5.7.6
	org.apache.avro	trevni-core	1.7.6-cdh5.7.6
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.7.6
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.7.6
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.7.6
	org.kitesdk	kite-data-hive	1.0.0-cdh5.7.6
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.7.6
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.7.6
	org.kitesdk	kite-data-s3	1.0.0-cdh5.7.6
	org.kitesdk	kite-data-spark	1.0.0-cdh5.7.6
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.7.6
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.7.6
	org.kitesdk	kite-minicluster	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.7.6
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.7.6
	org.kitesdk	kite-tools	1.0.0-cdh5.7.6
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.7.6

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-classification	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-core	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-demo	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-facet	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-join	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-memory	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-misc	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-queries	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.7.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.7.6
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.7.6
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.7.6
	org.apache.solr	solr-cell	4.10.3-cdh5.7.6
	org.apache.solr	solr-clustering	4.10.3-cdh5.7.6
	org.apache.solr	solr-core	4.10.3-cdh5.7.6
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.7.6
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.7.6
	org.apache.solr	solr-depends-sentry-lib	4.10.3-cdh5.7.6
	org.apache.solr	solr-langid	4.10.3-cdh5.7.6
	org.apache.solr	solr-security-util	4.10.3-cdh5.7.6
	org.apache.solr	solr-solrj	4.10.3-cdh5.7.6
	org.apache.solr	solr-test-framework	4.10.3-cdh5.7.6
	org.apache.solr	solr-uima	4.10.3-cdh5.7.6
	org.apache.solr	solr-velocity	4.10.3-cdh5.7.6
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.7.6
	com.cloudera.search	search-mr	1.0.0-cdh5.7.6
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.7.6
	com.ngdata	hbase-indexer-cli	1.5-cdh5.7.6

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-common	1.5-cdh5.7.6
	com.ngdata	hbase-indexer-demo	1.5-cdh5.7.6
	com.ngdata	hbase-indexer-dist	1.5-cdh5.7.6
	com.ngdata	hbase-indexer-engine	1.5-cdh5.7.6
	com.ngdata	hbase-indexer-model	1.5-cdh5.7.6
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.7.6
	com.ngdata	hbase-indexer-mr	1.5-cdh5.7.6
	com.ngdata	hbase-indexer-server	1.5-cdh5.7.6
	com.ngdata	hbase-sep-api	1.5-cdh5.7.6
	com.ngdata	hbase-sep-demo	1.5-cdh5.7.6
	com.ngdata	hbase-sep-impl	1.5-hbase1. 2-cdh5.7.6
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.7.6
	com.ngdata	hbase-sep-tools	1.5-cdh5.7.6

CDH 5.7.5

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.7.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.7.5

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.7.5
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.7.5
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.7.5
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.7.5
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.7.5
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.7.5
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.7.5
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.7.5
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.7.5
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.7.5

Version and Download Information

Project	groupId	artifactId	version
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.7.5
	org.apache.hive	hive-ant	1.1.0-cdh5.7.5
	org.apache.hive	hive-beeline	1.1.0-cdh5.7.5
	org.apache.hive	hive-cli	1.1.0-cdh5.7.5
	org.apache.hive	hive-common	1.1.0-cdh5.7.5
	org.apache.hive	hive-contrib	1.1.0-cdh5.7.5
	org.apache.hive	hive-exec	1.1.0-cdh5.7.5
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.7.5
	org.apache.hive	hive-hwi	1.1.0-cdh5.7.5
	org.apache.hive	hive-jdbc	1.1.0-cdh5.7.5
	org.apache.hive	hive-metastore	1.1.0-cdh5.7.5
	org.apache.hive	hive-serde	1.1.0-cdh5.7.5
	org.apache.hive	hive-service	1.1.0-cdh5.7.5
	org.apache.hive	hive-shims	1.1.0-cdh5.7.5
	org.apache.hive	hive-testutils	1.1.0-cdh5.7.5
	org.apache.hive	spark-client	1.1.0-cdh5.7.5
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.7.5
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.7.5
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.7.5
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.7.5
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.7.5

Project	groupId	artifactId	version
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.7.5
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.7.5
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.7.5
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.7.5
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-client	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-common	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-examples	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-it	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-rest	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-server	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.7.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shell	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-spark	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.7.5
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.7.5
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.7.5
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.7.5
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.7.5
	org.apache.pig	piggybank	0.12.0-cdh5.7.5
	org.apache.pig	pigsmoke	0.12.0-cdh5.7.5
	org.apache.pig	pigunit	0.12.0-cdh5.7.5
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.7.5
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.7.5
	org.apache.flume	flume-ng-core	1.6.0-cdh5.7.5
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.7.5
	org.apache.flume	flume-ng-node	1.6.0-cdh5.7.5
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.7.5
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.7.5
	org.apache.flume	flume-tools	1.6.0-cdh5.7.5
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.7.5
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.7.5

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.7.5
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.7.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.7.5
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.7.5
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.7.5
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-core	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-examples	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.7.5.oozie-4.1.0-cdh5.7.5
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.7.5.oozie-4.1.0-cdh5.7.5
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.7.5.oozie-4.1.0-cdh5.7.5
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.7.5.oozie-4.1.0-cdh5.7.5
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.7.5.oozie-4.1.0-cdh5.7.5
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.7.5.oozie-4.1.0-cdh5.7.5
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.7.5

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.7.5
	org.apache.oozie	oozie-tools	4.1.0-cdh5.7.5
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.7.5
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.7.5
	org.apache.mahout	mahout-core	0.9-cdh5.7.5
	org.apache.mahout	mahout-examples	0.9-cdh5.7.5
	org.apache.mahout	mahout-integration	0.9-cdh5.7.5
	org.apache.mahout	mahout-math	0.9-cdh5.7.5
	org.apache.mahout	mahout-math-scala	0.9-cdh5.7.5
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-chef	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-cli	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-core	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-druid	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.7.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-examples	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-hama	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-pig	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-solr	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.7.5
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.7.5
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.7.5
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.7.5
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.7.5
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.7.5
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.7.5
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.7.5
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.7.5
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.7.5
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.7.5

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.7.5
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.7.5
	org.apache.sqoop	test	1.99.5-cdh5.7.5
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.7.5
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.7.5
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.7.5
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.7.5
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.7.5
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.7.5
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.7.5
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.7.5
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.7.5
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.7.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-dist	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.7.5
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.7.5
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.7.5
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.7.5
	com.twitter	parquet-cascading	1.5.0-cdh5.7.5
	com.twitter	parquet-column	1.5.0-cdh5.7.5
	com.twitter	parquet-common	1.5.0-cdh5.7.5
	com.twitter	parquet-encoding	1.5.0-cdh5.7.5
	com.twitter	parquet-format	2.1.0-cdh5.7.5
	com.twitter	parquet-generator	1.5.0-cdh5.7.5

Project	groupId	artifactId	version
	com.twitter	parquet-hadoop	1.5.0-cdh5.7.5
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.7.5
	com.twitter	parquet-jackson	1.5.0-cdh5.7.5
	com.twitter	parquet-pig	1.5.0-cdh5.7.5
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.7.5
	com.twitter	parquet-protobuf	1.5.0-cdh5.7.5
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.7.5
	com.twitter	parquet-scooage_2.10	1.5.0-cdh5.7.5
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.7.5
	com.twitter	parquet-thrift	1.5.0-cdh5.7.5
	com.twitter	parquet-tools	1.5.0-cdh5.7.5
Llama	com.cloudera. llama	llama	1.0.0-cdh5.7.5
Apache Spark	org.apache. spark	spark-assembly_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-bagel_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-catalyst_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-core_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-graphx_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-hive_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-launcher_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-mllib_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-network-common_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-network-shuffle_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-network-yarn_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-repl_2.10	1.6.0-cdh5.7.5
	org.apache. spark	spark-sql_2.10	1.6.0-cdh5.7.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.7.5
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.7.5
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.7.5
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.7.5
	org.apache.crunch	crunch-core	0.11.0-cdh5.7.5
	org.apache.crunch	crunch-examples	0.11.0-cdh5.7.5
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.7.5
	org.apache.crunch	crunch-hive	0.11.0-cdh5.7.5
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.7.5
	org.apache.crunch	crunch-spark	0.11.0-cdh5.7.5
	org.apache.crunch	crunch-test	0.11.0-cdh5.7.5
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.7.5

Project	groupId	artifactId	version
	org.apache.avro	avro-compiler	1.7.6-cdh5.7.5
	org.apache.avro	avro-ipc	1.7.6-cdh5.7.5
	org.apache.avro	avro-mapred	1.7.6-cdh5.7.5
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.7.5
	org.apache.avro	avro-protobuf	1.7.6-cdh5.7.5
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.7.5
	org.apache.avro	avro-thrift	1.7.6-cdh5.7.5
	org.apache.avro	avro-tools	1.7.6-cdh5.7.5
	org.apache.avro	trevni-avro	1.7.6-cdh5.7.5
	org.apache.avro	trevni-core	1.7.6-cdh5.7.5
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.7.5
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.7.5
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.7.5
	org.kitesdk	kite-data-hive	1.0.0-cdh5.7.5
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.7.5
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.7.5
	org.kitesdk	kite-data-s3	1.0.0-cdh5.7.5
	org.kitesdk	kite-data-spark	1.0.0-cdh5.7.5
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.7.5
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.7.5
	org.kitesdk	kite-minicluster	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.7.5

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.7.5
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.7.5
	org.kitesdk	kite-tools	1.0.0-cdh5.7.5
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-classification	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-core	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-demo	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-facet	4.10.3-cdh5.7.5

Project	groupId	artifactId	version
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-join	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-memory	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-misc	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-queries	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.7.5
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.7.5
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.7.5
	org.apache.solr	solr-cell	4.10.3-cdh5.7.5
	org.apache.solr	solr-clustering	4.10.3-cdh5.7.5
	org.apache.solr	solr-core	4.10.3-cdh5.7.5
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.7.5
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.7.5
	org.apache.solr	solr-depends-sentry-libs	4.10.3-cdh5.7.5
	org.apache.solr	solr-langid	4.10.3-cdh5.7.5
	org.apache.solr	solr-security-util	4.10.3-cdh5.7.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-solrj	4.10.3-cdh5.7.5
	org.apache.solr	solr-test-framework	4.10.3-cdh5.7.5
	org.apache.solr	solr-uima	4.10.3-cdh5.7.5
	org.apache.solr	solr-velocity	4.10.3-cdh5.7.5
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.7.5
	com.cloudera.search	search-mr	1.0.0-cdh5.7.5
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.7.5
	com.ngdata	hbase-indexer-cli	1.5-cdh5.7.5
	com.ngdata	hbase-indexer-common	1.5-cdh5.7.5
	com.ngdata	hbase-indexer-demo	1.5-cdh5.7.5
	com.ngdata	hbase-indexer-dist	1.5-cdh5.7.5
	com.ngdata	hbase-indexer-engine	1.5-cdh5.7.5
	com.ngdata	hbase-indexer-model	1.5-cdh5.7.5
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.7.5
	com.ngdata	hbase-indexer-mr	1.5-cdh5.7.5
	com.ngdata	hbase-indexer-server	1.5-cdh5.7.5
	com.ngdata	hbase-sep-api	1.5-cdh5.7.5
	com.ngdata	hbase-sep-demo	1.5-cdh5.7.5
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.7.5
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.7.5
	com.ngdata	hbase-sep-tools	1.5-cdh5.7.5

CDH 5.7.4

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.7.4

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.7.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.7.4
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.7.4
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.7.4
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.7.4
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.7.4

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.7.4
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.7.4
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.7.4
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.7.4
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.7.4
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.7.4
	org.apache.hive	hive-ant	1.1.0-cdh5.7.4
	org.apache.hive	hive-beeline	1.1.0-cdh5.7.4
	org.apache.hive	hive-cli	1.1.0-cdh5.7.4
	org.apache.hive	hive-common	1.1.0-cdh5.7.4
	org.apache.hive	hive-contrib	1.1.0-cdh5.7.4
	org.apache.hive	hive-exec	1.1.0-cdh5.7.4
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.7.4
	org.apache.hive	hive-hwi	1.1.0-cdh5.7.4
	org.apache.hive	hive-jdbc	1.1.0-cdh5.7.4
	org.apache.hive	hive-metastore	1.1.0-cdh5.7.4
	org.apache.hive	hive-serde	1.1.0-cdh5.7.4
	org.apache.hive	hive-service	1.1.0-cdh5.7.4
	org.apache.hive	hive-shims	1.1.0-cdh5.7.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-testutils	1.1.0-cdh5.7.4
	org.apache.hive	spark-client	1.1.0-cdh5.7.4
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.7.4
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.7.4
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.7.4
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.7.4
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.7.4
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.7.4
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.7.4
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.7.4
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.7.4
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-client	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-common	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-examples	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-it	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.7.4

Project	groupId	artifactId	version
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-rest	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-server	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-shell	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-spark	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.7.4
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.7.4
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.7.4
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.7.4
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.7.4
	org.apache.pig	piggybank	0.12.0-cdh5.7.4
	org.apache.pig	pigsmoke	0.12.0-cdh5.7.4
	org.apache.pig	pigunit	0.12.0-cdh5.7.4
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.7.4
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.7.4
	org.apache.flume	flume-ng-core	1.6.0-cdh5.7.4
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.7.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-node	1.6.0-cdh5.7.4
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.7.4
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.7.4
	org.apache.flume	flume-tools	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.7.4
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.7.4

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.7.4
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.7.4
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.7.4
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.7.4
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.7.4
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.7.4
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.7.4
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.7.4
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.7.4
	org.apache. oozie	oozie-core	4.1.0-cdh5.7.4
	org.apache. oozie	oozie-examples	4.1.0-cdh5.7.4
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.7. 4.oozie-4.1. 0-cdh5.7.4
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.7. 4.oozie-4.1. 0-cdh5.7.4
	org.apache. oozie	oozie-hadoop-test	2.6. 0-mr1-cdh5.7. 4.oozie-4.1. 0-cdh5.7.4
	org.apache. oozie	oozie-hadoop-utils	2.6. 0-mr1-cdh5.7. 4.oozie-4.1. 0-cdh5.7.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.7.4.oozie-4.1.0-cdh5.7.4
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.7.4.oozie-4.1.0-cdh5.7.4
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.7.4
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.7.4
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.7.4
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.7.4
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.7.4
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.7.4
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.7.4
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.7.4
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.7.4
	org.apache.oozie	oozie-tools	4.1.0-cdh5.7.4
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.7.4
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.7.4
	org.apache.mahout	mahout-core	0.9-cdh5.7.4
	org.apache.mahout	mahout-examples	0.9-cdh5.7.4
	org.apache.mahout	mahout-integration	0.9-cdh5.7.4
	org.apache.mahout	mahout-math	0.9-cdh5.7.4
	org.apache.mahout	mahout-math-scala	0.9-cdh5.7.4
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.7.4

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-chef	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-cli	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-core	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-druid	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-examples	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-hama	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-pig	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-solr	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.7.4
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.7.4
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.7.4
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.7.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.7.4
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.7.4
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.7.4
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.7.4
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.7.4
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.7.4
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.7.4
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.7.4
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.7.4
	org.apache.sqoop	test	1.99.5-cdh5.7.4
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.7.4
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.7.4
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.7.4
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.7.4
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.7.4
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.7.4
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.7.4
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.7.4

Project	groupId	artifactId	version
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.7.4
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-dist	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.7.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.7.4
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.7.4
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.7.4
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.7.4
	com.twitter	parquet-cascading	1.5.0-cdh5.7.4
	com.twitter	parquet-column	1.5.0-cdh5.7.4
	com.twitter	parquet-common	1.5.0-cdh5.7.4
	com.twitter	parquet-encoding	1.5.0-cdh5.7.4
	com.twitter	parquet-format	2.1.0-cdh5.7.4
	com.twitter	parquet-generator	1.5.0-cdh5.7.4
	com.twitter	parquet-hadoop	1.5.0-cdh5.7.4
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.7.4
	com.twitter	parquet-jackson	1.5.0-cdh5.7.4
	com.twitter	parquet-pig	1.5.0-cdh5.7.4
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.7.4
	com.twitter	parquet-protobuf	1.5.0-cdh5.7.4
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.7.4
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.7.4
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.7.4
	com.twitter	parquet-thrift	1.5.0-cdh5.7.4
	com.twitter	parquet-tools	1.5.0-cdh5.7.4
Llama	com.cloudera.llama	llama	1.0.0-cdh5.7.4
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.7.4

Project	groupId	artifactId	version
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.7.4
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.7.4
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.7.4
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.7.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-core	0.11.0-cdh5.7.4
	org.apache.crunch	crunch-examples	0.11.0-cdh5.7.4
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.7.4
	org.apache.crunch	crunch-hive	0.11.0-cdh5.7.4
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.7.4
	org.apache.crunch	crunch-spark	0.11.0-cdh5.7.4
	org.apache.crunch	crunch-test	0.11.0-cdh5.7.4
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.7.4
	org.apache.avro	avro-compiler	1.7.6-cdh5.7.4
	org.apache.avro	avro-ipc	1.7.6-cdh5.7.4
	org.apache.avro	avro-mapred	1.7.6-cdh5.7.4
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.7.4
	org.apache.avro	avro-protobuf	1.7.6-cdh5.7.4
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.7.4
	org.apache.avro	avro-thrift	1.7.6-cdh5.7.4
	org.apache.avro	avro-tools	1.7.6-cdh5.7.4
	org.apache.avro	trevni-avro	1.7.6-cdh5.7.4
	org.apache.avro	trevni-core	1.7.6-cdh5.7.4
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.7.4
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.7.4
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.7.4
	org.kitesdk	kite-data-hive	1.0.0-cdh5.7.4
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.7.4
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.7.4

Project	groupId	artifactId	version
	org.kitesdk	kite-data-s3	1.0.0-cdh5.7.4
	org.kitesdk	kite-data-spark	1.0.0-cdh5.7.4
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.7.4
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.7.4
	org.kitesdk	kite-minicluster	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.7.4
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.7.4
	org.kitesdk	kite-tools	1.0.0-cdh5.7.4
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.7.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-classification	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-core	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-demo	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-facet	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-join	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-memory	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-misc	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-queries	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.7.4
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.7.4
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.7.4

Project	groupId	artifactId	version
	org.apache.solr	solr-cell	4.10.3-cdh5.7.4
	org.apache.solr	solr-clustering	4.10.3-cdh5.7.4
	org.apache.solr	solr-core	4.10.3-cdh5.7.4
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.7.4
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.7.4
	org.apache.solr	solr-depends-sentry-libs	4.10.3-cdh5.7.4
	org.apache.solr	solr-langid	4.10.3-cdh5.7.4
	org.apache.solr	solr-security-util	4.10.3-cdh5.7.4
	org.apache.solr	solr-solrj	4.10.3-cdh5.7.4
	org.apache.solr	solr-test-framework	4.10.3-cdh5.7.4
	org.apache.solr	solr-uima	4.10.3-cdh5.7.4
	org.apache.solr	solr-velocity	4.10.3-cdh5.7.4
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.7.4
	com.cloudera.search	search-mr	1.0.0-cdh5.7.4
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.7.4
	com.ngdata	hbase-indexer-cli	1.5-cdh5.7.4
	com.ngdata	hbase-indexer-common	1.5-cdh5.7.4
	com.ngdata	hbase-indexer-demo	1.5-cdh5.7.4
	com.ngdata	hbase-indexer-dist	1.5-cdh5.7.4
	com.ngdata	hbase-indexer-engine	1.5-cdh5.7.4
	com.ngdata	hbase-indexer-model	1.5-cdh5.7.4
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.7.4
	com.ngdata	hbase-indexer-mr	1.5-cdh5.7.4
	com.ngdata	hbase-indexer-server	1.5-cdh5.7.4
	com.ngdata	hbase-sep-api	1.5-cdh5.7.4
	com.ngdata	hbase-sep-demo	1.5-cdh5.7.4

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.7.4
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.7.4
	com.ngdata	hbase-sep-tools	1.5-cdh5.7.4

CDH 5.7.3

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.7.3

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.7.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.7.3
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.7.3
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.7.3
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.7.3
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.7.3
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.7.3
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.7.3
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.7.3
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.7.3
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.7.3
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.7.3
	org.apache.hive	hive-ant	1.1.0-cdh5.7.3
	org.apache.hive	hive-beeline	1.1.0-cdh5.7.3
	org.apache.hive	hive-cli	1.1.0-cdh5.7.3
	org.apache.hive	hive-common	1.1.0-cdh5.7.3
	org.apache.hive	hive-contrib	1.1.0-cdh5.7.3

Project	groupId	artifactId	version
	org.apache.hive	hive-exec	1.1.0-cdh5.7.3
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.7.3
	org.apache.hive	hive-hwi	1.1.0-cdh5.7.3
	org.apache.hive	hive-jdbc	1.1.0-cdh5.7.3
	org.apache.hive	hive-metastore	1.1.0-cdh5.7.3
	org.apache.hive	hive-serde	1.1.0-cdh5.7.3
	org.apache.hive	hive-service	1.1.0-cdh5.7.3
	org.apache.hive	hive-shims	1.1.0-cdh5.7.3
	org.apache.hive	hive-testutils	1.1.0-cdh5.7.3
	org.apache.hive	spark-client	1.1.0-cdh5.7.3
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.7.3
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.7.3
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.7.3
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.7.3
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.7.3
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.7.3
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.7.3
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.7.3
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.7.3
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.7.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-client	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-common	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-examples	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-it	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-rest	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-server	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-shell	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-spark	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.7.3
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.7.3
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.7.3
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.7.3

Project	groupId	artifactId	version
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.7.3
	org.apache.pig	piggybank	0.12.0-cdh5.7.3
	org.apache.pig	pigsmoke	0.12.0-cdh5.7.3
	org.apache.pig	pigunit	0.12.0-cdh5.7.3
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.7.3
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.7.3
	org.apache.flume	flume-ng-core	1.6.0-cdh5.7.3
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.7.3
	org.apache.flume	flume-ng-node	1.6.0-cdh5.7.3
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.7.3
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.7.3
	org.apache.flume	flume-tools	1.6.0-cdh5.7.3
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.7.3
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.7.3
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.7.3
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.7.3
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.7.3
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.7.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.7.3
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.7.3
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.7.3
	org.apache. oozie	oozie-core	4.1.0-cdh5.7.3

Project	groupId	artifactId	version
	org.apache.oozie	oozie-examples	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.7.3.oozie-4.1.0-cdh5.7.3
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.7.3.oozie-4.1.0-cdh5.7.3
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.7.3.oozie-4.1.0-cdh5.7.3
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.7.3.oozie-4.1.0-cdh5.7.3
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.7.3.oozie-4.1.0-cdh5.7.3
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.7.3.oozie-4.1.0-cdh5.7.3
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.7.3
	org.apache.oozie	oozie-tools	4.1.0-cdh5.7.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.7.3
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.7.3
	org.apache.mahout	mahout-core	0.9-cdh5.7.3
	org.apache.mahout	mahout-examples	0.9-cdh5.7.3
	org.apache.mahout	mahout-integration	0.9-cdh5.7.3
	org.apache.mahout	mahout-math	0.9-cdh5.7.3
	org.apache.mahout	mahout-math-scala	0.9-cdh5.7.3
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-chef	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-cli	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-core	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-druid	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-examples	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-hama	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.7.3

Project	groupId	artifactId	version
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-pig	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-solr	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.7.3
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.7.3
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.7.3
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.7.3
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.7.3
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.7.3
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.7.3
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.7.3
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.7.3
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.7.3
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.7.3
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.7.3
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.7.3
	org.apache.sqoop	test	1.99.5-cdh5.7.3
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.7.3
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.7.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.7.3
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.7.3
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.7.3
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.7.3
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.7.3
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.7.3
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.7.3
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-dist	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.7.3

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.7.3
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.7.3
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.7.3
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.7.3
	com.twitter	parquet-cascading	1.5.0-cdh5.7.3
	com.twitter	parquet-column	1.5.0-cdh5.7.3
	com.twitter	parquet-common	1.5.0-cdh5.7.3
	com.twitter	parquet-encoding	1.5.0-cdh5.7.3
	com.twitter	parquet-format	2.1.0-cdh5.7.3
	com.twitter	parquet-generator	1.5.0-cdh5.7.3
	com.twitter	parquet-hadoop	1.5.0-cdh5.7.3
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.7.3
	com.twitter	parquet-jackson	1.5.0-cdh5.7.3
	com.twitter	parquet-pig	1.5.0-cdh5.7.3
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.7.3
	com.twitter	parquet-protobuf	1.5.0-cdh5.7.3
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.7.3
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.7.3

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.7.3
	com.twitter	parquet-thrift	1.5.0-cdh5.7.3
	com.twitter	parquet-tools	1.5.0-cdh5.7.3
Llama	com.cloudera. llama	llama	1.0.0-cdh5.7.3
Apache Spark	org.apache. spark	spark-assembly_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-bagel_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-catalyst_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-core_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-graphx_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-hive_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-launcher_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-mllib_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-network-common_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-network-shuffle_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-network-yarn_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-repl_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-sql_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-streaming-flume_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-streaming-kafka_2.10	1.6.0-cdh5.7.3
	org.apache. spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.7.3

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.7.3
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.7.3
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.7.3
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.7.3
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.7.3
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.7.3
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.7.3
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.7.3
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.7.3
	org.apache.crunch	crunch-core	0.11.0-cdh5.7.3
	org.apache.crunch	crunch-examples	0.11.0-cdh5.7.3
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.7.3
	org.apache.crunch	crunch-hive	0.11.0-cdh5.7.3
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.7.3
	org.apache.crunch	crunch-spark	0.11.0-cdh5.7.3
	org.apache.crunch	crunch-test	0.11.0-cdh5.7.3
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.7.3
	org.apache.avro	avro-compiler	1.7.6-cdh5.7.3
	org.apache.avro	avro-ipc	1.7.6-cdh5.7.3
	org.apache.avro	avro-mapred	1.7.6-cdh5.7.3
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.7.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	avro-protobuf	1.7.6-cdh5.7.3
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.7.3
	org.apache.avro	avro-thrift	1.7.6-cdh5.7.3
	org.apache.avro	avro-tools	1.7.6-cdh5.7.3
	org.apache.avro	trevni-avro	1.7.6-cdh5.7.3
	org.apache.avro	trevni-core	1.7.6-cdh5.7.3
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.7.3
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.7.3
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.7.3
	org.kitesdk	kite-data-hive	1.0.0-cdh5.7.3
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.7.3
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.7.3
	org.kitesdk	kite-data-s3	1.0.0-cdh5.7.3
	org.kitesdk	kite-data-spark	1.0.0-cdh5.7.3
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.7.3
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.7.3
	org.kitesdk	kite-minicluster	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.7.3

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.7.3
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.7.3
	org.kitesdk	kite-tools	1.0.0-cdh5.7.3
Apache Solr	org.apache. lucene	lucene-analyzers-common	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-analyzers-icu	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-analyzers-morfologik	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-analyzers-phonetic	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-analyzers-smartcn	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-analyzers-stempel	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-analyzers-uima	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-benchmark	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-classification	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-codecs	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-core	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-demo	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-expressions	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-facet	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-grouping	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-highlighter	4.10.3-cdh5.7. 3
	org.apache. lucene	lucene-join	4.10.3-cdh5.7. 3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-memory	4.10.3-cdh5.7.3
	org.apache.lucene	lucene-misc	4.10.3-cdh5.7.3
	org.apache.lucene	lucene-queries	4.10.3-cdh5.7.3
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.7.3
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.7.3
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.7.3
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.7.3
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.7.3
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.7.3
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.7.3
	org.apache.solr	solr-cell	4.10.3-cdh5.7.3
	org.apache.solr	solr-clustering	4.10.3-cdh5.7.3
	org.apache.solr	solr-core	4.10.3-cdh5.7.3
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.7.3
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.7.3
	org.apache.solr	solr-depends-sentry-libs	4.10.3-cdh5.7.3
	org.apache.solr	solr-langid	4.10.3-cdh5.7.3
	org.apache.solr	solr-security-util	4.10.3-cdh5.7.3
	org.apache.solr	solr-solrj	4.10.3-cdh5.7.3
	org.apache.solr	solr-test-framework	4.10.3-cdh5.7.3
	org.apache.solr	solr-uima	4.10.3-cdh5.7.3

Project	groupId	artifactId	version
	org.apache.solr	solr-velocity	4.10.3-cdh5.7.3
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.7.3
	com.cloudera.search	search-mr	1.0.0-cdh5.7.3
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.7.3
	com.ngdata	hbase-indexer-cli	1.5-cdh5.7.3
	com.ngdata	hbase-indexer-common	1.5-cdh5.7.3
	com.ngdata	hbase-indexer-demo	1.5-cdh5.7.3
	com.ngdata	hbase-indexer-dist	1.5-cdh5.7.3
	com.ngdata	hbase-indexer-engine	1.5-cdh5.7.3
	com.ngdata	hbase-indexer-model	1.5-cdh5.7.3
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.7.3
	com.ngdata	hbase-indexer-mr	1.5-cdh5.7.3
	com.ngdata	hbase-indexer-server	1.5-cdh5.7.3
	com.ngdata	hbase-sep-api	1.5-cdh5.7.3
	com.ngdata	hbase-sep-demo	1.5-cdh5.7.3
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.7.3
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.7.3
	com.ngdata	hbase-sep-tools	1.5-cdh5.7.3

CDH 5.7.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.7.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.7.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.7.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.7.2
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.7.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.7.2
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.7.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.7.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.7.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.7.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.7.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.7.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.7.2
	org.apache.hive	hive-ant	1.1.0-cdh5.7.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.7.2
	org.apache.hive	hive-cli	1.1.0-cdh5.7.2
	org.apache.hive	hive-common	1.1.0-cdh5.7.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.7.2
	org.apache.hive	hive-exec	1.1.0-cdh5.7.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.7.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.7.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.7.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.7.2
	org.apache.hive	hive-serde	1.1.0-cdh5.7.2
	org.apache.hive	hive-service	1.1.0-cdh5.7.2
	org.apache.hive	hive-shims	1.1.0-cdh5.7.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.7.2
	org.apache.hive	spark-client	1.1.0-cdh5.7.2
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.7.2

Project	groupId	artifactId	version
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.7.2
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.7.2
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.7.2
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.7.2
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.7.2
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.7.2
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.7.2
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.7.2
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-client	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-common	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-examples	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-it	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.7.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-rest	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-server	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-shell	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-spark	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.7.2
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.7.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.7.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.7.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.7.2
	org.apache.pig	piggybank	0.12.0-cdh5.7.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.7.2
	org.apache.pig	pigunit	0.12.0-cdh5.7.2
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.7.2
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.7.2
	org.apache.flume	flume-ng-core	1.6.0-cdh5.7.2
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.7.2
	org.apache.flume	flume-ng-node	1.6.0-cdh5.7.2
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.7.2
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.7.2

Project	groupId	artifactId	version
	org.apache.flume	flume-tools	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.7.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sources	flume-jms-source	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.7.2
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.7.2
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-core	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-examples	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.7.2.oozie-4.1.0-cdh5.7.2
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.7.2.oozie-4.1.0-cdh5.7.2
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.7.2.oozie-4.1.0-cdh5.7.2
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.7.2.oozie-4.1.0-cdh5.7.2
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.7.2.oozie-4.1.0-cdh5.7.2
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.7.2.oozie-4.1.0-cdh5.7.2

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.7.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.7.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.7.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.7.2
	org.apache.mahout	mahout-core	0.9-cdh5.7.2
	org.apache.mahout	mahout-examples	0.9-cdh5.7.2
	org.apache.mahout	mahout-integration	0.9-cdh5.7.2
	org.apache.mahout	mahout-math	0.9-cdh5.7.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.7.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.7.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cli	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.7.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.7.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.7.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.7.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.7.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.7.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.7.2

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.7.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.7.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.7.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.7.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.7.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.7.2
	org.apache.sqoop	test	1.99.5-cdh5.7.2
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.7.2
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.7.2
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.7.2
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.7.2
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.7.2
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.7.2
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.7.2
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.7.2
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.7.2
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.7.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.7.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.7.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.7.2

Project	groupId	artifactId	version
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.7.2
	com.twitter	parquet-cascading	1.5.0-cdh5.7.2
	com.twitter	parquet-column	1.5.0-cdh5.7.2
	com.twitter	parquet-common	1.5.0-cdh5.7.2
	com.twitter	parquet-encoding	1.5.0-cdh5.7.2
	com.twitter	parquet-format	2.1.0-cdh5.7.2
	com.twitter	parquet-generator	1.5.0-cdh5.7.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.7.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.7.2
	com.twitter	parquet-jackson	1.5.0-cdh5.7.2
	com.twitter	parquet-pig	1.5.0-cdh5.7.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.7.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.7.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.7.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.7.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.7.2
	com.twitter	parquet-thrift	1.5.0-cdh5.7.2
com.twitter	parquet-tools	1.5.0-cdh5.7.2	
Llama	com.cloudera.llama	llama	1.0.0-cdh5.7.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.7.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.7.2
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.7.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.7.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.7.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.7.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.7.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.7.2

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hive	0.11.0-cdh5.7.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.7.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.7.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.7.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.7.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.7.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.7.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.7.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.7.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.7.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.7.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.7.2
	org.apache.avro	avro-tools	1.7.6-cdh5.7.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.7.2
	org.apache.avro	trevni-core	1.7.6-cdh5.7.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.7.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.7.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.7.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.7.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.7.2
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.7.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.7.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.7.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.7.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.7.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.7.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.7.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.7.2
	org.kitesdk	kite-tools	1.0.0-cdh5.7.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.7.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.7.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.7.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.7.2
	org.apache.solr	solr-cell	4.10.3-cdh5.7.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.7.2
	org.apache.solr	solr-core	4.10.3-cdh5.7.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.7.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.7.2
	org.apache.solr	solr-depends-sentry-libs	4.10.3-cdh5.7.2
	org.apache.solr	solr-langid	4.10.3-cdh5.7.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.7.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.7.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.7.2
	org.apache.solr	solr-uima	4.10.3-cdh5.7.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.7.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.7.2
	com.cloudera.search	search-mr	1.0.0-cdh5.7.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.7.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.7.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.7.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.7.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.7.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.7.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.7.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.7.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.7.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.7.2
	com.ngdata	hbase-sep-api	1.5-cdh5.7.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.7.2
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.7.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.7.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.7.2

CDH 5.7.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.7.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-mapreduce-client-native-task	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.7.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.7.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.7.1
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.7.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.7.1
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.7.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.7.1
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.7.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.7.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.7.1
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.7.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.7.1
	org.apache.hive	hive-ant	1.1.0-cdh5.7.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.7.1
	org.apache.hive	hive-cli	1.1.0-cdh5.7.1
	org.apache.hive	hive-common	1.1.0-cdh5.7.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.7.1
	org.apache.hive	hive-exec	1.1.0-cdh5.7.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.7.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.7.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-jdbc	1.1.0-cdh5.7.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.7.1
	org.apache.hive	hive-serde	1.1.0-cdh5.7.1
	org.apache.hive	hive-service	1.1.0-cdh5.7.1
	org.apache.hive	hive-shims	1.1.0-cdh5.7.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.7.1
	org.apache.hive	spark-client	1.1.0-cdh5.7.1
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.7.1
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.7.1
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.7.1
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.7.1
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.7.1
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.7.1
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.7.1
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.7.1
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.7.1
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-client	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-common	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-examples	1.2.0-cdh5.7.1

Project	groupId	artifactId	version
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-it	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-rest	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-server	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-shell	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-spark	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.7.1
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.7.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.7.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.7.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.7.1
	org.apache.pig	piggybank	0.12.0-cdh5.7.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.7.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.pig	pigunit	0.12.0-cdh5.7.1
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.7.1
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.7.1
	org.apache.flume	flume-ng-core	1.6.0-cdh5.7.1
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.7.1
	org.apache.flume	flume-ng-node	1.6.0-cdh5.7.1
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.7.1
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.7.1
	org.apache.flume	flume-tools	1.6.0-cdh5.7.1
	org.apache.flume. flume-ng-channels	flume-file-channel	1.6.0-cdh5.7.1
	org.apache.flume. flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.7.1
	org.apache.flume. flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.7.1
	org.apache.flume. flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.7.1
	org.apache.flume. flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.7.1
	org.apache.flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.7.1
	org.apache.flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.7.1
	org.apache.flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.7.1
	org.apache.flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.7.1

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.7.1
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.7.1
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.7.1
	org.apache. oozie	oozie-core	4.1.0-cdh5.7.1
	org.apache. oozie	oozie-examples	4.1.0-cdh5.7.1
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.7. 1.oozie-4.1. 0-cdh5.7.1
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.7.

Version and Download Information

Project	groupId	artifactId	version
			1.oozie-4.1.0-cdh5.7.1
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.7.1.oozie-4.1.0-cdh5.7.1
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.7.1.oozie-4.1.0-cdh5.7.1
	org.apache.oozie	oozie-hbase	1.2.0-cdh5.7.1.oozie-4.1.0-cdh5.7.1
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.7.1.oozie-4.1.0-cdh5.7.1
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.7.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.7.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.7.1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.7.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.7.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.7.1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.7.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.7.1
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.7.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.7.1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.7.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.7.1
	org.apache.mahout	mahout-core	0.9-cdh5.7.1
	org.apache.mahout	mahout-examples	0.9-cdh5.7.1

Project	groupId	artifactId	version
	org.apache.mahout	mahout-integration	0.9-cdh5.7.1
	org.apache.mahout	mahout-math	0.9-cdh5.7.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.7.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.7.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.7.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.7.1
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.7.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.7.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.7.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.7.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.7.1
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.7.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.7.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.7.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.7.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.7.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.7.1
	org.apache.sqoop	test	1.99.5-cdh5.7.1
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.7.1
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.7.1
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.7.1
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.7.1
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.7.1

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.7.1
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.7.1
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.7.1
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.7.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.7.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.7.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.7.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.7.1
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.7.1
	com.twitter	parquet-cascading	1.5.0-cdh5.7.1
	com.twitter	parquet-column	1.5.0-cdh5.7.1
	com.twitter	parquet-common	1.5.0-cdh5.7.1
	com.twitter	parquet-encoding	1.5.0-cdh5.7.1
	com.twitter	parquet-format	2.1.0-cdh5.7.1
	com.twitter	parquet-generator	1.5.0-cdh5.7.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.7.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.7.1
	com.twitter	parquet-jackson	1.5.0-cdh5.7.1
	com.twitter	parquet-pig	1.5.0-cdh5.7.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.7.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.7.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.7.1
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.7.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.7.1
	com.twitter	parquet-thrift	1.5.0-cdh5.7.1
	com.twitter	parquet-tools	1.5.0-cdh5.7.1
Llama	com.cloudera.llama	llama	1.0.0-cdh5.7.1
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.7.1

Project	groupId	artifactId	version
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.7.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.7.1
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.7.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.7.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.7.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.7.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.7.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.7.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.7.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.7.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.7.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.7.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.7.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.7.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.7.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.7.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.7.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.7.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.7.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.7.1
	org.apache.avro	avro-tools	1.7.6-cdh5.7.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.7.1

Project	groupId	artifactId	version
	org.apache.avro	trevni-core	1.7.6-cdh5.7.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.7.1
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.7.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.7.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.7.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.7.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.7.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.7.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.7.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.7.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.7.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.7.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.7.1
	org.kitesdk	kite-tools	1.0.0-cdh5.7.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.7.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-core	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-join	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.7.1

Project	groupId	artifactId	version
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.7.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.7.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.7.1
	org.apache.solr	solr-cell	4.10.3-cdh5.7.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.7.1
	org.apache.solr	solr-core	4.10.3-cdh5.7.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.7.1
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.7.1
	org.apache.solr	solr-depends-sentry-libs	4.10.3-cdh5.7.1
	org.apache.solr	solr-langid	4.10.3-cdh5.7.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.7.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.7.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.7.1
	org.apache.solr	solr-uima	4.10.3-cdh5.7.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.7.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.7.1
	com.cloudera.search	search-mr	1.0.0-cdh5.7.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.7.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.7.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.7.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.7.1

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-dist	1.5-cdh5.7.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.7.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.7.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.7.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.7.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.7.1
	com.ngdata	hbase-sep-api	1.5-cdh5.7.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.7.1
	com.ngdata	hbase-sep-impl	1.5-hbase1.2-cdh5.7.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.7.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.7.1

CDH 5.7.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.7.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.7.0
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.7.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.7.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.7.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.7.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.7.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.7.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.7.0
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.7.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.7.0
	org.apache.hive	hive-ant	1.1.0-cdh5.7.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.7.0
	org.apache.hive	hive-cli	1.1.0-cdh5.7.0
	org.apache.hive	hive-common	1.1.0-cdh5.7.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.7.0
	org.apache.hive	hive-exec	1.1.0-cdh5.7.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.7.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.7.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.7.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.7.0
	org.apache.hive	hive-serde	1.1.0-cdh5.7.0
	org.apache.hive	hive-service	1.1.0-cdh5.7.0
	org.apache.hive	hive-shims	1.1.0-cdh5.7.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.7.0
	org.apache.hive	spark-client	1.1.0-cdh5.7.0
	org.apache.hive	hive-hcatalog	1.1.0-cdh5.7.0
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.7.0
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.7.0
	org.apache.hive	hive-hcatalog-server-extensions	1.1.0-cdh5.7.0
	org.apache.hive	hive-hcatalog-streaming	1.1.0-cdh5.7.0
	org.apache.hive	hive-webhcat	1.1.0-cdh5.7.0
	org.apache.hive	hive-webhcat-java-client	1.1.0-cdh5.7.0
	org.apache.hive	hive-shims-0.23	1.1.0-cdh5.7.0
	org.apache.hive	hive-shims-common	1.1.0-cdh5.7.0
	org.apache.hive	hive-shims-scheduler	1.1.0-cdh5.7.0
Apache HBase	org.apache.hbase	hbase-annotations	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-checkstyle	1.2.0-cdh5.7.0

Project	groupId	artifactId	version
	org.apache.hbase	hbase-client	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-common	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-examples	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-external-blockcache	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-hadoop-compat	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-hadoop2-compat	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-it	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-prefix-tree	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-procedure	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-protocol	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-resource-bundle	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-rest	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-server	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-shaded-client	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-shaded-server	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-shell	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-spark	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-testing-util	1.2.0-cdh5.7.0
	org.apache.hbase	hbase-thrift	1.2.0-cdh5.7.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.7.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.7.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.7.0
	org.apache.pig	piggybank	0.12.0-cdh5.7.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.7.0
	org.apache.pig	pigunit	0.12.0-cdh5.7.0
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-core	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.7.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-node	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.7.0
	org.apache.flume	flume-tools	1.6.0-cdh5.7.0
	org.apache.flume	flume-file-channel	1.6.0-cdh5.7.0
	org.apache.flume	flume-jdbc-channel	1.6.0-cdh5.7.0
	org.apache.flume	flume-kafka-channel	1.6.0-cdh5.7.0
	org.apache.flume	flume-spillable-memory-channel	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-log4jappender	1.6.0-cdh5.7.0
	org.apache.flume	flume-avro-source	1.6.0-cdh5.7.0
	org.apache.flume	flume-thrift-source	1.6.0-cdh5.7.0
	org.apache.flume	flume-dataset-sink	1.6.0-cdh5.7.0
	org.apache.flume	flume-hdfs-sink	1.6.0-cdh5.7.0
	org.apache.flume	flume-hive-sink	1.6.0-cdh5.7.0
	org.apache.flume	flume-irc-sink	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-elasticsearch-sink	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-hbase-sink	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-kafka-sink	1.6.0-cdh5.7.0
	org.apache.flume	flume-ng-morphline-solr-sink	1.6.0-cdh5.7.0
	org.apache.flume	flume-jms-source	1.6.0-cdh5.7.0
	org.apache.flume	flume-kafka-source	1.6.0-cdh5.7.0
	org.apache.flume	flume-scribe-source	1.6.0-cdh5.7.0
	org.apache.flume	flume-taildir-source	1.6.0-cdh5.7.0
	org.apache.flume	flume-twitter-source	1.6.0-cdh5.7.0
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-core	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-examples	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.7.0

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.7.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.7.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.7.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.7.0
	org.apache.mahout	mahout-core	0.9-cdh5.7.0
	org.apache.mahout	mahout-examples	0.9-cdh5.7.0
	org.apache.mahout	mahout-integration	0.9-cdh5.7.0
	org.apache.mahout	mahout-math	0.9-cdh5.7.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.7.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.7.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.7.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.7.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.7.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.7.0
	org.apache.sqoop	test	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.7.0
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.7.0

Project	groupId	artifactId	version
	org.apache.sentry	sentry-submission	1.99.5-cdh5.7.0
Apache Sentry	org.apache.sentry	sentry-submission-mapreduce	1.99.5-cdh5.7.0
	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.7.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.7.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.7.0
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.7.0
	com.twitter	parquet-cascading	1.5.0-cdh5.7.0
	com.twitter	parquet-column	1.5.0-cdh5.7.0
	com.twitter	parquet-common	1.5.0-cdh5.7.0
	com.twitter	parquet-encoding	1.5.0-cdh5.7.0
	com.twitter	parquet-format	2.1.0-cdh5.7.0

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-generator	1.5.0-cdh5.7.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.7.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.7.0
	com.twitter	parquet-jackson	1.5.0-cdh5.7.0
	com.twitter	parquet-pig	1.5.0-cdh5.7.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.7.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.7.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.7.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.7.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.7.0
	com.twitter	parquet-thrift	1.5.0-cdh5.7.0
	com.twitter	parquet-tools	1.5.0-cdh5.7.0
Llama	com.doudera.llama	llama	1.0.0-cdh5.7.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-bagel_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-catalyst_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-core_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-docker-integration-tests_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-graphx_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-hive_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-launcher_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-mllib_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-network-common_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-network-shuffle_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-network-yarn_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-repl_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-sql_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-streaming-flume_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-streaming-kafka_2.10	1.6.0-cdh5.7.0

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-streaming-twitter_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-streaming_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-test-tags_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-unsafe_2.10	1.6.0-cdh5.7.0
	org.apache.spark	spark-yarn_2.10	1.6.0-cdh5.7.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.7.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.7.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.7.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.7.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.7.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.7.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.7.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.7.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.7.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.7.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.7.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.7.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.7.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.7.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.7.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.7.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.7.0
	org.apache.avro	avro-tools	1.7.6-cdh5.7.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.7.0
	org.apache.avro	trevni-core	1.7.6-cdh5.7.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.7.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.7.0

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.7.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.7.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.7.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.7.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.7.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.7.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.7.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.7.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-metrics-scalable	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.7.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.7.0
	org.kitesdk	kite-tools	1.0.0-cdh5.7.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.7.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.7.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.7.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.7.0
	org.apache.solr	solr-cell	4.10.3-cdh5.7.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.7.0
	org.apache.solr	solr-core	4.10.3-cdh5.7.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.7.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.7.0
	org.apache.solr	solr-langid	4.10.3-cdh5.7.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.7.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.7.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.7.0
	org.apache.solr	solr-uima	4.10.3-cdh5.7.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.7.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.7.0
	com.cloudera.search	search-mr	1.0.0-cdh5.7.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.7.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.7.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.7.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.7.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.7.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.7.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.7.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.7.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.7.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.7.0
	com.ngdata	hbase-sep-api	1.5-cdh5.7.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.7.0
	com.ngdata	hbase-sep-impl	1.5-cdh5.7.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.7.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.7.0

Maven Artifacts for CDH 5.6.x Releases

CDH 5.6.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.6.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.6.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.6.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.6.1
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.6.1

Project	groupId	artifactId	version
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.6.1
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.6.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.6.1
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.6.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.6.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.6.1
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.6.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.6.1
	org.apache.hive	hive-ant	1.1.0-cdh5.6.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.6.1
	org.apache.hive	hive-cli	1.1.0-cdh5.6.1
	org.apache.hive	hive-common	1.1.0-cdh5.6.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.6.1
	org.apache.hive	hive-exec	1.1.0-cdh5.6.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.6.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.6.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.6.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.6.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-serde	1.1.0-cdh5.6.1
	org.apache.hive	hive-service	1.1.0-cdh5.6.1
	org.apache.hive	hive-shims	1.1.0-cdh5.6.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.6.1
	org.apache.hive	spark-client	1.1.0-cdh5.6.1
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.6.1
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.6.1
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.6.1
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.6.1
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.6.1
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.6.1
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.6.1
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.6.1
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.6.1
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-client	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-common	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-examples	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.6.1

Project	groupId	artifactId	version
	org.apache.hbase	hbase-it	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-rest	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-server	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-shell	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.6.1
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.6.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.6.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.6.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.6.1
	org.apache.pig	piggybank	0.12.0-cdh5.6.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.6.1
	org.apache.pig	pigunit	0.12.0-cdh5.6.1
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.6.1
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.6.1
	org.apache.flume	flume-ng-core	1.6.0-cdh5.6.1
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.6.1
	org.apache.flume	flume-ng-node	1.6.0-cdh5.6.1
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.6.1
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.6.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-tools	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.6.1
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.6.1

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.6.1
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.6.1
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.6.1
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.6.1
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.6.1
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.6.1
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.6.1
	org.apache. oozie	oozie-core	4.1.0-cdh5.6.1
	org.apache. oozie	oozie-examples	4.1.0-cdh5.6.1
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.6. 1.oozie-4.1. 0-cdh5.6.1
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.6. 1.oozie-4.1. 0-cdh5.6.1
	org.apache. oozie	oozie-hadoop-test	2.6. 0-mr1-cdh5.6. 1.oozie-4.1. 0-cdh5.6.1
	org.apache. oozie	oozie-hadoop-utils	2.6. 0-mr1-cdh5.6. 1.oozie-4.1. 0-cdh5.6.1
	org.apache. oozie	oozie-hbase	1.0.0-cdh5.6.1. oozie-4.1. 0-cdh5.6.1
	org.apache. oozie	oozie-hcatalog	1.1.0-cdh5.6.1. oozie-4.1. 0-cdh5.6.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.6.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.6.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.6.1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.6.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.6.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.6.1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.6.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.6.1
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.6.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.6.1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.6.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.6.1
	org.apache.mahout	mahout-core	0.9-cdh5.6.1
	org.apache.mahout	mahout-examples	0.9-cdh5.6.1
	org.apache.mahout	mahout-integration	0.9-cdh5.6.1
	org.apache.mahout	mahout-math	0.9-cdh5.6.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.6.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.6.1

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cli	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.6.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.6.1
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.6.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.6.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.6.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.6.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.6.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.6.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.6.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.6.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.6.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.6.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.6.1
	org.apache.sqoop	test	1.99.5-cdh5.6.1
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.6.1
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.6.1
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.6.1
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.6.1
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.6.1
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.6.1
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.6.1
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.6.1
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.6.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.6.1

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.6.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.6.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.6.1

Version and Download Information

Project	groupId	artifactId	version	
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.6.1	
	com.twitter	parquet-cascading	1.5.0-cdh5.6.1	
	com.twitter	parquet-column	1.5.0-cdh5.6.1	
	com.twitter	parquet-common	1.5.0-cdh5.6.1	
	com.twitter	parquet-encoding	1.5.0-cdh5.6.1	
	com.twitter	parquet-format	2.1.0-cdh5.6.1	
	com.twitter	parquet-generator	1.5.0-cdh5.6.1	
	com.twitter	parquet-hadoop	1.5.0-cdh5.6.1	
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.6.1	
	com.twitter	parquet-jackson	1.5.0-cdh5.6.1	
	com.twitter	parquet-pig	1.5.0-cdh5.6.1	
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.6.1	
	com.twitter	parquet-protobuf	1.5.0-cdh5.6.1	
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.6.1	
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.6.1	
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.6.1	
	com.twitter	parquet-thrift	1.5.0-cdh5.6.1	
	com.twitter	parquet-tools	1.5.0-cdh5.6.1	
	Llama	com.cloudera.llama	llama	1.0.0-cdh5.6.1
	Apache Spark	org.apache.spark	spark-assembly_2.10	1.5.0-cdh5.6.1
com.databricks		spark-avro_2.10	1.1.0-cdh5.6.1	
org.apache.spark		spark-bagel_2.10	1.5.0-cdh5.6.1	
org.apache.spark		spark-catalyst_2.10	1.5.0-cdh5.6.1	
org.apache.spark		spark-core_2.10	1.5.0-cdh5.6.1	
org.apache.spark		spark-graphx_2.10	1.5.0-cdh5.6.1	
org.apache.spark		spark-hive_2.10	1.5.0-cdh5.6.1	
org.apache.spark		spark-launcher_2.10	1.5.0-cdh5.6.1	
org.apache.spark		spark-mllib_2.10	1.5.0-cdh5.6.1	

Project	groupId	artifactId	version
	org.apache.spark	spark-network-common_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-network-shuffle_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-network-yarn_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-repl_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-sql_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-streaming-flume_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-streaming-kafka_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-streaming-twitter_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-streaming_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-unsafe_2.10	1.5.0-cdh5.6.1
	org.apache.spark	spark-yarn_2.10	1.5.0-cdh5.6.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.6.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.6.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.6.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.6.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.6.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.6.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.6.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.6.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.6.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.6.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.6.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.6.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.6.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.6.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.6.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.6.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.6.1
	org.apache.avro	avro-tools	1.7.6-cdh5.6.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.6.1
	org.apache.avro	trevni-core	1.7.6-cdh5.6.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.6.1
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.6.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.6.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.6.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.6.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.6.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.6.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.6.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.6.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.6.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.6.1

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.6.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.6.1
	org.kitesdk	kite-tools	1.0.0-cdh5.6.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.6.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-core	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-join	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.6.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.6.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.6.1
	org.apache.solr	solr-cell	4.10.3-cdh5.6.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.6.1
	org.apache.solr	solr-core	4.10.3-cdh5.6.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.6.1

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.6.1
	org.apache.solr	solr-langid	4.10.3-cdh5.6.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.6.1
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.6.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.6.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.6.1
	org.apache.solr	solr-uima	4.10.3-cdh5.6.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.6.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.6.1
	com.cloudera.search	search-mr	1.0.0-cdh5.6.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.6.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.6.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.6.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.6.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.6.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.6.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.6.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.6.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.6.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.6.1
	com.ngdata	hbase-sep-api	1.5-cdh5.6.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.6.1
	com.ngdata	hbase-sep-impl	1.5-hbase1.0-cdh5.6.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.6.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.6.1

CDH 5.6.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Version and Download Information

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.6.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.6.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-core	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-examples	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-streaming	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-test	2.6.0-cdh5.6.0
	org.apache.hadoop	hadoop-tools	2.6.0-cdh5.6.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.6.0
	org.apache.hive	hive-ant	1.1.0-cdh5.6.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.6.0
	org.apache.hive	hive-cli	1.1.0-cdh5.6.0
	org.apache.hive	hive-common	1.1.0-cdh5.6.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.6.0
	org.apache.hive	hive-exec	1.1.0-cdh5.6.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.6.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.6.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.6.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-metastore	1.1.0-cdh5.6.0
	org.apache.hive	hive-serde	1.1.0-cdh5.6.0
	org.apache.hive	hive-service	1.1.0-cdh5.6.0
	org.apache.hive	hive-shims	1.1.0-cdh5.6.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.6.0
	org.apache.hive	spark-client	1.1.0-cdh5.6.0
	org.apache.hcatalog	hive-hcatalog-core	1.1.0-cdh5.6.0
	org.apache.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.6.0
	org.apache.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.6.0
	org.apache.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.6.0
	org.apache.hcatalog	hive-webhcat	1.1.0-cdh5.6.0
	org.apache.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.6.0
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.6.0
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.6.0
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.6.0
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-client	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-common	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-examples	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-it	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-rest	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-server	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-shell	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.6.0
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.6.0

Project	groupId	artifactId	version
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.6.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.6.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.6.0
	org.apache.pig	piggybank	0.12.0-cdh5.6.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.6.0
	org.apache.pig	pigunit	0.12.0-cdh5.6.0
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-core	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-node	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.6.0
	org.apache.flume	flume-tools	1.6.0-cdh5.6.0
	org.apache.flume	flume-file-channel	1.6.0-cdh5.6.0
	org.apache.flume	flume-jdbc-channel	1.6.0-cdh5.6.0
	org.apache.flume	flume-kafka-channel	1.6.0-cdh5.6.0
	org.apache.flume	flume-spillable-memory-channel	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-log4jappender	1.6.0-cdh5.6.0
	org.apache.flume	flume-avro-source	1.6.0-cdh5.6.0
	org.apache.flume	flume-thrift-source	1.6.0-cdh5.6.0
	org.apache.flume	flume-dataset-sink	1.6.0-cdh5.6.0
	org.apache.flume	flume-hdfs-sink	1.6.0-cdh5.6.0
	org.apache.flume	flume-hive-sink	1.6.0-cdh5.6.0
	org.apache.flume	flume-irc-sink	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-elasticsearch-sink	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-hbase-sink	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-kafka-sink	1.6.0-cdh5.6.0
	org.apache.flume	flume-ng-morphline-solr-sink	1.6.0-cdh5.6.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-jms-source	1.6.0-cdh5.6.0
	org.apache.flume	flume-kafka-source	1.6.0-cdh5.6.0
	org.apache.flume	flume-scribe-source	1.6.0-cdh5.6.0
	org.apache.flume	flume-taildir-source	1.6.0-cdh5.6.0
	org.apache.flume	flume-twitter-source	1.6.0-cdh5.6.0
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-core	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-examples	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-hadoop	2.0.1-cdh5.6.0
	org.apache.oozie	oozie-hadoop-distcp	2.0.1-cdh5.6.0
	org.apache.oozie	oozie-hadoop-test	2.0.1-cdh5.6.0
	org.apache.oozie	oozie-hadoop-utils	2.0.1-cdh5.6.0
	org.apache.oozie	oozie-hbase	1.0.0-cdh5.6.0
	org.apache.oozie	oozie-hcatalog	1.0.0-cdh5.6.0
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.6.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.6.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.6.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.6.0
	org.apache.mahout	mahout-core	0.9-cdh5.6.0
	org.apache.mahout	mahout-examples	0.9-cdh5.6.0
	org.apache.mahout	mahout-integration	0.9-cdh5.6.0
	org.apache.mahout	mahout-math	0.9-cdh5.6.0

Project	groupId	artifactId	version
	org.apache.mahout	mahout-math-scala	0.9-cdh5.6.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.6.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.6.0
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.6.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.6.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.6.0
	org.apache.sqoop	test	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.6.0
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.6.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.6.0

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.6.0
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.6.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.6.0
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.6.0
	com.twitter	parquet-cascading	1.5.0-cdh5.6.0
	com.twitter	parquet-column	1.5.0-cdh5.6.0
	com.twitter	parquet-common	1.5.0-cdh5.6.0
	com.twitter	parquet-encoding	1.5.0-cdh5.6.0
	com.twitter	parquet-format	2.1.0-cdh5.6.0
	com.twitter	parquet-generator	1.5.0-cdh5.6.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.6.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.6.0
	com.twitter	parquet-jackson	1.5.0-cdh5.6.0
	com.twitter	parquet-pig	1.5.0-cdh5.6.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.6.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.6.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.6.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.6.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.6.0
	com.twitter	parquet-thrift	1.5.0-cdh5.6.0
	com.twitter	parquet-tools	1.5.0-cdh5.6.0
Llama	com.doudera.llama	llama	1.0.0-cdh5.6.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-bagel_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-catalyst_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-core_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-graphx_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-hive_2.10	1.5.0-cdh5.6.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-launcher_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-mllib_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-network-common_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-network-shuffle_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-network-yarn_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-repl_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-sql_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-streaming-flume_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-streaming-kafka_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-streaming-twitter_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-streaming_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-unsafe_2.10	1.5.0-cdh5.6.0
	org.apache.spark	spark-yarn_2.10	1.5.0-cdh5.6.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.6.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.6.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.6.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.6.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.6.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.6.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.6.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.6.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.6.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.6.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.6.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.6.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.6.0

Project	groupId	artifactId	version
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.6.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.6.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.6.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.6.0
	org.apache.avro	avro-tools	1.7.6-cdh5.6.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.6.0
	org.apache.avro	trevni-core	1.7.6-cdh5.6.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.6.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.6.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.6.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.6.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.6.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.6.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.6.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.6.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.6.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.6.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.6.0

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.6.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.6.0
	org.kitesdk	kite-tools	1.0.0-cdh5.6.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.6.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.6.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.6.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.6.0
	org.apache.solr	solr-cell	4.10.3-cdh5.6.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.6.0
	org.apache.solr	solr-core	4.10.3-cdh5.6.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.6.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.6.0
	org.apache.solr	solr-langid	4.10.3-cdh5.6.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.6.0
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.6.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.6.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.6.0
	org.apache.solr	solr-uima	4.10.3-cdh5.6.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.6.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.6.0
	com.cloudera.search	search-mr	1.0.0-cdh5.6.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.6.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.6.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.6.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.6.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.6.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.6.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.6.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.6.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.6.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.6.0
	com.ngdata	hbase-sep-api	1.5-cdh5.6.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.6.0

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-impl	1.5-cdh5.6.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.6.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.6.0

Maven Artifacts for CDH 5.5.x Releases CDH 5.5.6

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.5.6

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.5.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.5.6
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.5.6
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.5.6
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.5.6
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.5.6
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.5.6
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.5.6
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.5.6
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.5.6
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.5.6
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.5.6
	org.apache.hive	hive-ant	1.1.0-cdh5.5.6
	org.apache.hive	hive-beeline	1.1.0-cdh5.5.6
	org.apache.hive	hive-cli	1.1.0-cdh5.5.6
	org.apache.hive	hive-common	1.1.0-cdh5.5.6

Project	groupId	artifactId	version
	org.apache.hive	hive-contrib	1.1.0-cdh5.5.6
	org.apache.hive	hive-exec	1.1.0-cdh5.5.6
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.5.6
	org.apache.hive	hive-hwi	1.1.0-cdh5.5.6
	org.apache.hive	hive-jdbc	1.1.0-cdh5.5.6
	org.apache.hive	hive-metastore	1.1.0-cdh5.5.6
	org.apache.hive	hive-serde	1.1.0-cdh5.5.6
	org.apache.hive	hive-service	1.1.0-cdh5.5.6
	org.apache.hive	hive-shims	1.1.0-cdh5.5.6
	org.apache.hive	hive-testutils	1.1.0-cdh5.5.6
	org.apache.hive	spark-client	1.1.0-cdh5.5.6
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.5.6
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.5.6
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.5.6
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.5.6
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.5.6
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.5.6
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.5.6
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.5.6
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.5.6
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.5.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-client	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-common	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-examples	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-it	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-rest	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-server	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-shell	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.5.6
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.5.6
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.5.6
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.5.6
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.5.6
	org.apache.pig	piggybank	0.12.0-cdh5.5.6
	org.apache.pig	pigsmoke	0.12.0-cdh5.5.6
	org.apache.pig	pigunit	0.12.0-cdh5.5.6
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.5.6

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.5.6
	org.apache.flume	flume-ng-core	1.6.0-cdh5.5.6
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.5.6
	org.apache.flume	flume-ng-node	1.6.0-cdh5.5.6
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.5.6
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.5.6
	org.apache.flume	flume-tools	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-channels	flume-file-channel	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sources	flume-avro-source	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.5.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sources	flume-jms-source	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sources	flume-taildir-source	1.6.0-cdh5.5.6
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.6.0-cdh5.5.6
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-core	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-examples	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.5.6.oozie-4.1.0-cdh5.5.6
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.5.6.oozie-4.1.0-cdh5.5.6
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.5.

Project	groupId	artifactId	version
			6.oozie-4.1. 0-cdh5.5.6
	org.apache.oozie	oozie-hadoop-utils	2.6. 0-mr1-cdh5.5. 6.oozie-4.1. 0-cdh5.5.6
	org.apache.oozie	oozie-hbase	1.0.0-cdh5.5.6. oozie-4.1. 0-cdh5.5.6
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.5.6. oozie-4.1. 0-cdh5.5.6
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.5.6
	org.apache.oozie	oozie-tools	4.1.0-cdh5.5.6
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.5.6
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.5.6
	org.apache.mahout	mahout-core	0.9-cdh5.5.6
	org.apache.mahout	mahout-examples	0.9-cdh5.5.6
	org.apache.mahout	mahout-integration	0.9-cdh5.5.6
	org.apache.mahout	mahout-math	0.9-cdh5.5.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.mahout	mahout-math-scala	0.9-cdh5.5.6
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-chef	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-cli	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-core	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-druid	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-examples	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-hama	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-pig	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-solr	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.5.6
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.5.6

Project	groupId	artifactId	version
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.5.6
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.5.6
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.5.6
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.5.6
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.5.6
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.5.6
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.5.6
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.5.6
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.5.6
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.5.6
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.5.6
	org.apache.sqoop	test	1.99.5-cdh5.5.6
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.5.6
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.5.6
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.5.6
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.5.6
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.5.6
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.5.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.5.6
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.5.6
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.5.6
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-dist	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.5.6

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.5.6
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.5.6
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.5.6
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.5.6
	com.twitter	parquet-cascading	1.5.0-cdh5.5.6
	com.twitter	parquet-column	1.5.0-cdh5.5.6
	com.twitter	parquet-common	1.5.0-cdh5.5.6
	com.twitter	parquet-encoding	1.5.0-cdh5.5.6
	com.twitter	parquet-format	2.1.0-cdh5.5.6
	com.twitter	parquet-generator	1.5.0-cdh5.5.6
	com.twitter	parquet-hadoop	1.5.0-cdh5.5.6
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.5.6
	com.twitter	parquet-jackson	1.5.0-cdh5.5.6
	com.twitter	parquet-pig	1.5.0-cdh5.5.6
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.5.6
	com.twitter	parquet-protobuf	1.5.0-cdh5.5.6
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.5.6
	com.twitter	parquet-scooage_2.10	1.5.0-cdh5.5.6
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.5.6
	com.twitter	parquet-thrift	1.5.0-cdh5.5.6
	com.twitter	parquet-tools	1.5.0-cdh5.5.6
Llama	com.cloudera.llama	llama	1.0.0-cdh5.5.6
Apache Spark	org.apache.spark	spark-assembly_2.10	1.5.0-cdh5.5.6
	com.databricks	spark-avro_2.10	1.1.0-cdh5.5.6
	org.apache.spark	spark-bagel_2.10	1.5.0-cdh5.5.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-catalyst_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-core_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-graphx_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-hive_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-launcher_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-mllib_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-network-common_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-network-shuffle_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-network-yarn_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-repl_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-sql_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-streaming-flume-sink_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-streaming-flume_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-streaming-kafka_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-streaming-mqtt_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-streaming-twitter_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-streaming-zeromq_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-streaming_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-unsafe_2.10	1.5.0-cdh5.5.6
	org.apache.spark	spark-yarn_2.10	1.5.0-cdh5.5.6

Project	groupId	artifactId	version
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.5.6
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.5.6
	org.apache.crunch	crunch-core	0.11.0-cdh5.5.6
	org.apache.crunch	crunch-examples	0.11.0-cdh5.5.6
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.5.6
	org.apache.crunch	crunch-hive	0.11.0-cdh5.5.6
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.5.6
	org.apache.crunch	crunch-spark	0.11.0-cdh5.5.6
	org.apache.crunch	crunch-test	0.11.0-cdh5.5.6
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.5.6
	org.apache.avro	avro-compiler	1.7.6-cdh5.5.6
	org.apache.avro	avro-ipc	1.7.6-cdh5.5.6
	org.apache.avro	avro-mapred	1.7.6-cdh5.5.6
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.5.6
	org.apache.avro	avro-protobuf	1.7.6-cdh5.5.6
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.5.6
	org.apache.avro	avro-thrift	1.7.6-cdh5.5.6
	org.apache.avro	avro-tools	1.7.6-cdh5.5.6
	org.apache.avro	trevni-avro	1.7.6-cdh5.5.6
	org.apache.avro	trevni-core	1.7.6-cdh5.5.6
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.5.6
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.5.6

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.5.6
	org.kitesdk	kite-data-hive	1.0.0-cdh5.5.6
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.5.6
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.5.6
	org.kitesdk	kite-data-s3	1.0.0-cdh5.5.6
	org.kitesdk	kite-data-spark	1.0.0-cdh5.5.6
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.5.6
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.5.6
	org.kitesdk	kite-minicluster	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.5.6
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.5.6
	org.kitesdk	kite-tools	1.0.0-cdh5.5.6
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.5.6

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-classification	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-core	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-demo	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-facet	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-join	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-memory	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-misc	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-queries	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.5.6
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.5.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.5.6
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.5.6
	org.apache.solr	solr-cell	4.10.3-cdh5.5.6
	org.apache.solr	solr-clustering	4.10.3-cdh5.5.6
	org.apache.solr	solr-core	4.10.3-cdh5.5.6
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.5.6
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.5.6
	org.apache.solr	solr-langid	4.10.3-cdh5.5.6
	org.apache.solr	solr-security-util	4.10.3-cdh5.5.6
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.5.6
	org.apache.solr	solr-solrj	4.10.3-cdh5.5.6
	org.apache.solr	solr-test-framework	4.10.3-cdh5.5.6
	org.apache.solr	solr-uima	4.10.3-cdh5.5.6
	org.apache.solr	solr-velocity	4.10.3-cdh5.5.6
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.5.6
	com.cloudera.search	search-mr	1.0.0-cdh5.5.6
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.5.6
	com.ngdata	hbase-indexer-cli	1.5-cdh5.5.6
	com.ngdata	hbase-indexer-common	1.5-cdh5.5.6
	com.ngdata	hbase-indexer-demo	1.5-cdh5.5.6
	com.ngdata	hbase-indexer-dist	1.5-cdh5.5.6
	com.ngdata	hbase-indexer-engine	1.5-cdh5.5.6
	com.ngdata	hbase-indexer-model	1.5-cdh5.5.6
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.5.6
	com.ngdata	hbase-indexer-mr	1.5-cdh5.5.6

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-server	1.5-cdh5.5.6
	com.ngdata	hbase-sep-api	1.5-cdh5.5.6
	com.ngdata	hbase-sep-demo	1.5-cdh5.5.6
	com.ngdata	hbase-sep-impl	1.5-hbase1.0-cdh5.5.6
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.5.6
	com.ngdata	hbase-sep-tools	1.5-cdh5.5.6

CDH 5.5.5

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.5.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.5.5

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.5.5
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.5.5
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.5.5
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.5.5
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.5.5
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.5.5
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.5.5
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.5.5
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.5.5
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.5.5
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.5.5
	org.apache.hive	hive-ant	1.1.0-cdh5.5.5
	org.apache.hive	hive-beeline	1.1.0-cdh5.5.5
	org.apache.hive	hive-cli	1.1.0-cdh5.5.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-common	1.1.0-cdh5.5.5
	org.apache.hive	hive-contrib	1.1.0-cdh5.5.5
	org.apache.hive	hive-exec	1.1.0-cdh5.5.5
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.5.5
	org.apache.hive	hive-hwi	1.1.0-cdh5.5.5
	org.apache.hive	hive-jdbc	1.1.0-cdh5.5.5
	org.apache.hive	hive-metastore	1.1.0-cdh5.5.5
	org.apache.hive	hive-serde	1.1.0-cdh5.5.5
	org.apache.hive	hive-service	1.1.0-cdh5.5.5
	org.apache.hive	hive-shims	1.1.0-cdh5.5.5
	org.apache.hive	hive-testutils	1.1.0-cdh5.5.5
	org.apache.hive	spark-client	1.1.0-cdh5.5.5
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.5.5
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.5.5
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.5.5
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.5.5
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.5.5
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.5.5
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.5.5
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.5.5
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.5.5

Project	groupId	artifactId	version
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-client	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-common	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-examples	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-it	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-rest	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-server	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-shell	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.5.5
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.5.5
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.5.5
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.5.5
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.5.5
	org.apache.pig	piggybank	0.12.0-cdh5.5.5
	org.apache.pig	pigsmoke	0.12.0-cdh5.5.5
	org.apache.pig	pigunit	0.12.0-cdh5.5.5

Version and Download Information

Project	groupId	artifactId	version
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.5.5
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.5.5
	org.apache.flume	flume-ng-core	1.6.0-cdh5.5.5
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.5.5
	org.apache.flume	flume-ng-node	1.6.0-cdh5.5.5
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.5.5
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.5.5
	org.apache.flume	flume-tools	1.6.0-cdh5.5.5
	org.apache.flume. flume-ng-channels	flume-file-channel	1.6.0-cdh5.5.5
	org.apache.flume. flume-ng-channels	flume-jdbc-channel	1.6.0-cdh5.5.5
	org.apache.flume. flume-ng-channels	flume-kafka-channel	1.6.0-cdh5.5.5
	org.apache.flume. flume-ng-channels	flume-spillable-memory-channel	1.6.0-cdh5.5.5
	org.apache.flume. flume-ng-clients	flume-ng-log4jappender	1.6.0-cdh5.5.5
	org.apache.flume. flume-ng-sources	flume-avro-source	1.6.0-cdh5.5.5
	org.apache.flume. flume-ng-sources	flume-thrift-source	1.6.0-cdh5.5.5
	org.apache.flume. flume-ng-sinks	flume-dataset-sink	1.6.0-cdh5.5.5
	org.apache.flume. flume-ng-sinks	flume-hdfs-sink	1.6.0-cdh5.5.5

Project	groupId	artifactId	version
	org.apache. flume. flume-ng-sinks	flume-hive-sink	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sinks	flume-irc-sink	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sinks	flume-ng-elasticsearch-sink	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sinks	flume-ng-hbase-sink	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sinks	flume-ng-kafka-sink	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sinks	flume-ng-morphline-solr-sink	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sources	flume-jms-source	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sources	flume-kafka-source	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sources	flume-scribe-source	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sources	flume-taildir-source	1.6.0-cdh5.5.5
	org.apache. flume. flume-ng-sources	flume-twitter-source	1.6.0-cdh5.5.5
Apache Oozie	org.apache. oozie	oozie-client	4.1.0-cdh5.5.5
	org.apache. oozie	oozie-core	4.1.0-cdh5.5.5
	org.apache. oozie	oozie-examples	4.1.0-cdh5.5.5
	org.apache. oozie	oozie-hadoop	2.6. 0-mr1-cdh5.5. 5.oozie-4.1. 0-cdh5.5.5
	org.apache. oozie	oozie-hadoop-distcp	2.6. 0-mr1-cdh5.5.

Version and Download Information

Project	groupId	artifactId	version
			5.oozie-4.1.0-cdh5.5.5
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.5.5.oozie-4.1.0-cdh5.5.5
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.5.5.oozie-4.1.0-cdh5.5.5
	org.apache.oozie	oozie-hbase	1.0.0-cdh5.5.5.oozie-4.1.0-cdh5.5.5
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.5.5.oozie-4.1.0-cdh5.5.5
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.5.5
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.5.5
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.5.5
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.5.5
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.5.5
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.5.5
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.5.5
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.5.5
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.5.5
	org.apache.oozie	oozie-tools	4.1.0-cdh5.5.5
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.5.5
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.5.5
	org.apache.mahout	mahout-core	0.9-cdh5.5.5
	org.apache.mahout	mahout-examples	0.9-cdh5.5.5

Project	groupId	artifactId	version
	org.apache.mahout	mahout-integration	0.9-cdh5.5.5
	org.apache.mahout	mahout-math	0.9-cdh5.5.5
	org.apache.mahout	mahout-math-scala	0.9-cdh5.5.5
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-chef	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-cli	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-core	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-druid	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-examples	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-hama	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-pig	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-solr	0.9.0-cdh5.5.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.5.5
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.5.5
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.5.5
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.5.5
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.5.5
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.5.5
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.5.5
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.5.5
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.5.5
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.5.5
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.5.5
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.5.5
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.5.5
	org.apache.sqoop	test	1.99.5-cdh5.5.5
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.5.5
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.5.5
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.5.5
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.5.5
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.5.5

Project	groupId	artifactId	version
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.5.5
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.5.5
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.5.5
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.5.5
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-dist	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.5.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.5.5
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.5.5
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.5.5
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.5.5
	com.twitter	parquet-cascading	1.5.0-cdh5.5.5
	com.twitter	parquet-column	1.5.0-cdh5.5.5
	com.twitter	parquet-common	1.5.0-cdh5.5.5
	com.twitter	parquet-encoding	1.5.0-cdh5.5.5
	com.twitter	parquet-format	2.1.0-cdh5.5.5
	com.twitter	parquet-generator	1.5.0-cdh5.5.5
	com.twitter	parquet-hadoop	1.5.0-cdh5.5.5
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.5.5
	com.twitter	parquet-jackson	1.5.0-cdh5.5.5
	com.twitter	parquet-pig	1.5.0-cdh5.5.5
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.5.5
	com.twitter	parquet-protobuf	1.5.0-cdh5.5.5
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.5.5
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.5.5
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.5.5
	com.twitter	parquet-thrift	1.5.0-cdh5.5.5
	com.twitter	parquet-tools	1.5.0-cdh5.5.5
Llama	com.cloudera.llama	llama	1.0.0-cdh5.5.5
Apache Spark	org.apache.spark	spark-assembly_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-bagel_2.10	1.5.0-cdh5.5.5

Project	groupId	artifactId	version
	org.apache.spark	spark-catalyst_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-core_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-graphx_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-hive_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-launcher_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-mllib_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-network-common_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-network-shuffle_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-network-yarn_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-repl_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-sql_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-streaming-flume-sink_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-streaming-flume_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-streaming-kafka_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-streaming-mqtt_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-streaming-twitter_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-streaming-zeromq_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-streaming_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-unsafe_2.10	1.5.0-cdh5.5.5
	org.apache.spark	spark-yarn_2.10	1.5.0-cdh5.5.5

Version and Download Information

Project	groupId	artifactId	version
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.5.5
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.5.5
	org.apache.crunch	crunch-core	0.11.0-cdh5.5.5
	org.apache.crunch	crunch-examples	0.11.0-cdh5.5.5
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.5.5
	org.apache.crunch	crunch-hive	0.11.0-cdh5.5.5
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.5.5
	org.apache.crunch	crunch-spark	0.11.0-cdh5.5.5
	org.apache.crunch	crunch-test	0.11.0-cdh5.5.5
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.5.5
	org.apache.avro	avro-compiler	1.7.6-cdh5.5.5
	org.apache.avro	avro-ipc	1.7.6-cdh5.5.5
	org.apache.avro	avro-mapred	1.7.6-cdh5.5.5
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.5.5
	org.apache.avro	avro-protobuf	1.7.6-cdh5.5.5
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.5.5
	org.apache.avro	avro-thrift	1.7.6-cdh5.5.5
	org.apache.avro	avro-tools	1.7.6-cdh5.5.5
	org.apache.avro	trevni-avro	1.7.6-cdh5.5.5
	org.apache.avro	trevni-core	1.7.6-cdh5.5.5
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.5.5
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.5.5

Project	groupId	artifactId	version
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.5.5
	org.kitesdk	kite-data-hive	1.0.0-cdh5.5.5
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.5.5
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.5.5
	org.kitesdk	kite-data-s3	1.0.0-cdh5.5.5
	org.kitesdk	kite-data-spark	1.0.0-cdh5.5.5
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.5.5
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.5.5
	org.kitesdk	kite-minicluster	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.5.5
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.5.5
	org.kitesdk	kite-tools	1.0.0-cdh5.5.5
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.5.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-classification	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-core	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-demo	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-facet	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-join	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-memory	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-misc	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-queries	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.5.5
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.5.5

Project	groupId	artifactId	version
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.5.5
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.5.5
	org.apache.solr	solr-cell	4.10.3-cdh5.5.5
	org.apache.solr	solr-clustering	4.10.3-cdh5.5.5
	org.apache.solr	solr-core	4.10.3-cdh5.5.5
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.5.5
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.5.5
	org.apache.solr	solr-langid	4.10.3-cdh5.5.5
	org.apache.solr	solr-security-util	4.10.3-cdh5.5.5
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.5.5
	org.apache.solr	solr-solrj	4.10.3-cdh5.5.5
	org.apache.solr	solr-test-framework	4.10.3-cdh5.5.5
	org.apache.solr	solr-uima	4.10.3-cdh5.5.5
	org.apache.solr	solr-velocity	4.10.3-cdh5.5.5
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.5.5
	com.cloudera.search	search-mr	1.0.0-cdh5.5.5
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.5.5
	com.ngdata	hbase-indexer-cli	1.5-cdh5.5.5
	com.ngdata	hbase-indexer-common	1.5-cdh5.5.5
	com.ngdata	hbase-indexer-demo	1.5-cdh5.5.5
	com.ngdata	hbase-indexer-dist	1.5-cdh5.5.5
	com.ngdata	hbase-indexer-engine	1.5-cdh5.5.5
	com.ngdata	hbase-indexer-model	1.5-cdh5.5.5
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.5.5
	com.ngdata	hbase-indexer-mr	1.5-cdh5.5.5

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-server	1.5-cdh5.5.5
	com.ngdata	hbase-sep-api	1.5-cdh5.5.5
	com.ngdata	hbase-sep-demo	1.5-cdh5.5.5
	com.ngdata	hbase-sep-impl	1.5-hbase1.0-cdh5.5.5
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.5.5
	com.ngdata	hbase-sep-tools	1.5-cdh5.5.5

CDH 5.5.4

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.5.4

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.5.4
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.5.4
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.5.4
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.5.4
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.5.4
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.5.4
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.5.4
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.5.4
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.5.4
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.5.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-ant	1.1.0-cdh5.5.4
	org.apache.hive	hive-beeline	1.1.0-cdh5.5.4
	org.apache.hive	hive-cli	1.1.0-cdh5.5.4
	org.apache.hive	hive-common	1.1.0-cdh5.5.4
	org.apache.hive	hive-contrib	1.1.0-cdh5.5.4
	org.apache.hive	hive-exec	1.1.0-cdh5.5.4
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.5.4
	org.apache.hive	hive-hwi	1.1.0-cdh5.5.4
	org.apache.hive	hive-jdbc	1.1.0-cdh5.5.4
	org.apache.hive	hive-metastore	1.1.0-cdh5.5.4
	org.apache.hive	hive-serde	1.1.0-cdh5.5.4
	org.apache.hive	hive-service	1.1.0-cdh5.5.4
	org.apache.hive	hive-shims	1.1.0-cdh5.5.4
	org.apache.hive	hive-testutils	1.1.0-cdh5.5.4
	org.apache.hive	spark-client	1.1.0-cdh5.5.4
	org.apache.hcatalog	hive-hcatalog-core	1.1.0-cdh5.5.4
	org.apache.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.5.4
	org.apache.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.5.4
	org.apache.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.5.4
	org.apache.hcatalog	hive-webhcat	1.1.0-cdh5.5.4
	org.apache.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.5.4
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.5.4
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.5.4
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.5.4
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-client	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-common	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-examples	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.5.4

Project	groupId	artifactId	version
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-it	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-rest	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-server	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-shell	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.5.4
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.5.4
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.5.4
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.5.4
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.5.4
	org.apache.pig	piggybank	0.12.0-cdh5.5.4
	org.apache.pig	pigsmoke	0.12.0-cdh5.5.4
	org.apache.pig	pigunit	0.12.0-cdh5.5.4
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-core	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-node	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.5.4
	org.apache.flume	flume-tools	1.6.0-cdh5.5.4
	org.apache.flume	flume-file-channel	1.6.0-cdh5.5.4
	org.apache.flume	flume-jdbc-channel	1.6.0-cdh5.5.4
	org.apache.flume	flume-kafka-channel	1.6.0-cdh5.5.4
	org.apache.flume	flume-spillable-memory-channel	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-log4jappender	1.6.0-cdh5.5.4
	org.apache.flume	flume-avro-source	1.6.0-cdh5.5.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-thrift-source	1.6.0-cdh5.5.4
	org.apache.flume	flume-dataset-sink	1.6.0-cdh5.5.4
	org.apache.flume	flume-hdfs-sink	1.6.0-cdh5.5.4
	org.apache.flume	flume-hive-sink	1.6.0-cdh5.5.4
	org.apache.flume	flume-irc-sink	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-elasticsearch-sink	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-hbase-sink	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-kafka-sink	1.6.0-cdh5.5.4
	org.apache.flume	flume-ng-morphline-solr-sink	1.6.0-cdh5.5.4
	org.apache.flume	flume-jms-source	1.6.0-cdh5.5.4
	org.apache.flume	flume-kafka-source	1.6.0-cdh5.5.4
	org.apache.flume	flume-scribe-source	1.6.0-cdh5.5.4
	org.apache.flume	flume-taildir-source	1.6.0-cdh5.5.4
	org.apache.flume	flume-twitter-source	1.6.0-cdh5.5.4
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-core	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-examples	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.5.4

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.5.4
	org.apache.oozie	oozie-tools	4.1.0-cdh5.5.4
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.5.4
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.5.4
	org.apache.mahout	mahout-core	0.9-cdh5.5.4
	org.apache.mahout	mahout-examples	0.9-cdh5.5.4
	org.apache.mahout	mahout-integration	0.9-cdh5.5.4
	org.apache.mahout	mahout-math	0.9-cdh5.5.4
	org.apache.mahout	mahout-math-scala	0.9-cdh5.5.4
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-chef	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-cli	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-core	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-druid	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-examples	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-hama	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-pig	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-solr	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.5.4
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.5.4

Version and Download Information

Project	groupId	artifactId	version
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.5.4
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.5.4
	org.apache.sqoop	test	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.5.4
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.5.4
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-dist	1.5.1-cdh5.5.4
org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.5.4	

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.5.4
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.5.4
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.5.4
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.5.4
	com.twitter	parquet-cascading	1.5.0-cdh5.5.4
	com.twitter	parquet-column	1.5.0-cdh5.5.4
	com.twitter	parquet-common	1.5.0-cdh5.5.4
	com.twitter	parquet-encoding	1.5.0-cdh5.5.4
	com.twitter	parquet-format	2.1.0-cdh5.5.4
	com.twitter	parquet-generator	1.5.0-cdh5.5.4
	com.twitter	parquet-hadoop	1.5.0-cdh5.5.4
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.5.4
	com.twitter	parquet-jackson	1.5.0-cdh5.5.4
	com.twitter	parquet-pig	1.5.0-cdh5.5.4
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.5.4
	com.twitter	parquet-protobuf	1.5.0-cdh5.5.4
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.5.4
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.5.4
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.5.4

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-thrift	1.5.0-cdh5.5.4
	com.twitter	parquet-tools	1.5.0-cdh5.5.4
Llama	com.doudera.llama	llama	1.0.0-cdh5.5.4
Apache Spark	org.apache.spark	spark-assembly_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-bagel_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-catalyst_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-core_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-graphx_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-hive_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-launcher_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-mllib_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-network-common_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-network-shuffle_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-network-yarn_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-repl_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-sql_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-streaming-flume-sink_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-streaming-flume_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-streaming-kafka_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-streaming-mqtt_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-streaming-twitter_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-streaming-zeromq_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-streaming_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-unsafe_2.10	1.5.0-cdh5.5.4
	org.apache.spark	spark-yarn_2.10	1.5.0-cdh5.5.4
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.5.4
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.5.4
	org.apache.crunch	crunch-core	0.11.0-cdh5.5.4
	org.apache.crunch	crunch-examples	0.11.0-cdh5.5.4

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.5.4
	org.apache.crunch	crunch-hive	0.11.0-cdh5.5.4
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.5.4
	org.apache.crunch	crunch-spark	0.11.0-cdh5.5.4
	org.apache.crunch	crunch-test	0.11.0-cdh5.5.4
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.5.4
	org.apache.avro	avro-compiler	1.7.6-cdh5.5.4
	org.apache.avro	avro-ipc	1.7.6-cdh5.5.4
	org.apache.avro	avro-mapred	1.7.6-cdh5.5.4
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.5.4
	org.apache.avro	avro-protobuf	1.7.6-cdh5.5.4
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.5.4
	org.apache.avro	avro-thrift	1.7.6-cdh5.5.4
	org.apache.avro	avro-tools	1.7.6-cdh5.5.4
	org.apache.avro	trevni-avro	1.7.6-cdh5.5.4
	org.apache.avro	trevni-core	1.7.6-cdh5.5.4
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.5.4
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.5.4
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.5.4
	org.kitesdk	kite-data-hive	1.0.0-cdh5.5.4
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.5.4
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.5.4
	org.kitesdk	kite-data-s3	1.0.0-cdh5.5.4
	org.kitesdk	kite-data-spark	1.0.0-cdh5.5.4
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.5.4
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.5.4
	org.kitesdk	kite-minicluster	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.5.4

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.5.4
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.5.4
	org.kitesdk	kite-tools	1.0.0-cdh5.5.4
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-classification	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-core	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-demo	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-facet	4.10.3-cdh5.5.4

Project	groupId	artifactId	version
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-join	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-memory	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-misc	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-queries	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.5.4
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.5.4
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.5.4
	org.apache.solr	solr-cell	4.10.3-cdh5.5.4
	org.apache.solr	solr-clustering	4.10.3-cdh5.5.4
	org.apache.solr	solr-core	4.10.3-cdh5.5.4
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.5.4
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.5.4
	org.apache.solr	solr-langid	4.10.3-cdh5.5.4
	org.apache.solr	solr-security-util	4.10.3-cdh5.5.4
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.5.4
	org.apache.solr	solr-solrj	4.10.3-cdh5.5.4
	org.apache.solr	solr-test-framework	4.10.3-cdh5.5.4
	org.apache.solr	solr-uima	4.10.3-cdh5.5.4
	org.apache.solr	solr-velocity	4.10.3-cdh5.5.4
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.5.4
	com.cloudera.search	search-mr	1.0.0-cdh5.5.4
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.5.4
	com.ngdata	hbase-indexer-cli	1.5-cdh5.5.4
	com.ngdata	hbase-indexer-common	1.5-cdh5.5.4

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-demo	1.5-cdh5.5.4
	com.ngdata	hbase-indexer-dist	1.5-cdh5.5.4
	com.ngdata	hbase-indexer-engine	1.5-cdh5.5.4
	com.ngdata	hbase-indexer-model	1.5-cdh5.5.4
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.5.4
	com.ngdata	hbase-indexer-mr	1.5-cdh5.5.4
	com.ngdata	hbase-indexer-server	1.5-cdh5.5.4
	com.ngdata	hbase-sep-api	1.5-cdh5.5.4
	com.ngdata	hbase-sep-demo	1.5-cdh5.5.4
	com.ngdata	hbase-sep-impl	1.5-cdh5.5.4
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.5.4
	com.ngdata	hbase-sep-tools	1.5-cdh5.5.4

CDH 5.5.2

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.5.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.5.2
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.5.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.5.2
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.5.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.5.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.5.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.5.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.5.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.5.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.5.2
	org.apache.hive	hive-ant	1.1.0-cdh5.5.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.5.2
	org.apache.hive	hive-cli	1.1.0-cdh5.5.2
	org.apache.hive	hive-common	1.1.0-cdh5.5.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.5.2
	org.apache.hive	hive-exec	1.1.0-cdh5.5.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.5.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.5.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.5.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.5.2
	org.apache.hive	hive-serde	1.1.0-cdh5.5.2
	org.apache.hive	hive-service	1.1.0-cdh5.5.2
	org.apache.hive	hive-shims	1.1.0-cdh5.5.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.5.2
	org.apache.hive	spark-client	1.1.0-cdh5.5.2
	org.apache.hive	hive-hcatalog	1.1.0-cdh5.5.2
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.5.2
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.5.2
	org.apache.hive	hive-hcatalog-server-extensions	1.1.0-cdh5.5.2
	org.apache.hive	hive-hcatalog-streaming	1.1.0-cdh5.5.2
	org.apache.hive	hive-webhcat	1.1.0-cdh5.5.2
	org.apache.hive	hive-webhcat-java-client	1.1.0-cdh5.5.2
	org.apache.hive	hive-shims-0.23	1.1.0-cdh5.5.2
	org.apache.hive	hive-shims-common	1.1.0-cdh5.5.2
	org.apache.hive	hive-shims-scheduler	1.1.0-cdh5.5.2

Project	groupId	artifactId	version
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-client	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-common	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-examples	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-it	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-rest	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-server	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-shell	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.5.2
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.5.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.5.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.5.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.5.2
	org.apache.pig	piggybank	0.12.0-cdh5.5.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.5.2
	org.apache.pig	pigunit	0.12.0-cdh5.5.2
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-core	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-node	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.5.2
	org.apache.flume	flume-tools	1.6.0-cdh5.5.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-file-channel	1.6.0-cdh5.5.2
	org.apache.flume	flume-jdbc-channel	1.6.0-cdh5.5.2
	org.apache.flume	flume-kafka-channel	1.6.0-cdh5.5.2
	org.apache.flume	flume-spillable-memory-channel	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-log4jappender	1.6.0-cdh5.5.2
	org.apache.flume	flume-avro-source	1.6.0-cdh5.5.2
	org.apache.flume	flume-thrift-source	1.6.0-cdh5.5.2
	org.apache.flume	flume-dataset-sink	1.6.0-cdh5.5.2
	org.apache.flume	flume-hdfs-sink	1.6.0-cdh5.5.2
	org.apache.flume	flume-hive-sink	1.6.0-cdh5.5.2
	org.apache.flume	flume-irc-sink	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-elasticsearch-sink	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-hbase-sink	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-kafka-sink	1.6.0-cdh5.5.2
	org.apache.flume	flume-ng-morphline-solr-sink	1.6.0-cdh5.5.2
	org.apache.flume	flume-jms-source	1.6.0-cdh5.5.2
	org.apache.flume	flume-kafka-source	1.6.0-cdh5.5.2
	org.apache.flume	flume-scribe-source	1.6.0-cdh5.5.2
	org.apache.flume	flume-taildir-source	1.6.0-cdh5.5.2
	org.apache.flume	flume-twitter-source	1.6.0-cdh5.5.2
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-core	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-examples	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.5.2

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.5.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.5.2
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.5.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.5.2
	org.apache.mahout	mahout-core	0.9-cdh5.5.2
	org.apache.mahout	mahout-examples	0.9-cdh5.5.2
	org.apache.mahout	mahout-integration	0.9-cdh5.5.2
	org.apache.mahout	mahout-math	0.9-cdh5.5.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.5.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.5.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.5.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.5.2
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.5.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.5.2
	org.apache.sqoop	test	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.5.2
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.5.2
Apache Sentry (Incubating)	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.5.2

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-dist	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.5.2
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.5.2
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.5.2
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.5.2
	com.twitter	parquet-cascading	1.5.0-cdh5.5.2
	com.twitter	parquet-column	1.5.0-cdh5.5.2
	com.twitter	parquet-common	1.5.0-cdh5.5.2
	com.twitter	parquet-encoding	1.5.0-cdh5.5.2
	com.twitter	parquet-format	2.1.0-cdh5.5.2
	com.twitter	parquet-generator	1.5.0-cdh5.5.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.5.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.5.2

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-jackson	1.5.0-cdh5.5.2
	com.twitter	parquet-pig	1.5.0-cdh5.5.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.5.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.5.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.5.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.5.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.5.2
	com.twitter	parquet-thrift	1.5.0-cdh5.5.2
	com.twitter	parquet-tools	1.5.0-cdh5.5.2
Llama	com.doudera.llama	llama	1.0.0-cdh5.5.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-bagel_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-catalyst_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-core_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-graphx_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-hive_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-launcher_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-mllib_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-network-common_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-network-shuffle_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-network-yarn_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-repl_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-sql_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-streaming-flume_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-streaming-kafka_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-streaming-twitter_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.5.0-cdh5.5.2

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-unsafe_2.10	1.5.0-cdh5.5.2
	org.apache.spark	spark-yarn_2.10	1.5.0-cdh5.5.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.5.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.5.2
	org.apache.crunch	crunch-core	0.11.0-cdh5.5.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.5.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.5.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.5.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.5.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.5.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.5.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.5.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.5.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.5.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.5.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.5.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.5.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.5.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.5.2
	org.apache.avro	avro-tools	1.7.6-cdh5.5.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.5.2
	org.apache.avro	trevni-core	1.7.6-cdh5.5.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.5.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.5.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.5.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.5.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.5.2
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.5.2
	org.kitesdk	kite-data-s3	1.0.0-cdh5.5.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-spark	1.0.0-cdh5.5.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.5.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.5.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.5.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.5.2
	org.kitesdk	kite-tools	1.0.0-cdh5.5.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.5.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.5.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.5.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.5.2
	org.apache.solr	solr-cell	4.10.3-cdh5.5.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.5.2
	org.apache.solr	solr-core	4.10.3-cdh5.5.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.5.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.5.2
	org.apache.solr	solr-langid	4.10.3-cdh5.5.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.5.2
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.5.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.5.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.5.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-uima	4.10.3-cdh5.5.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.5.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.5.2
	com.cloudera.search	search-mr	1.0.0-cdh5.5.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.5.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.5.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.5.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.5.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.5.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.5.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.5.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.5.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.5.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.5.2
	com.ngdata	hbase-sep-api	1.5-cdh5.5.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.5.2
	com.ngdata	hbase-sep-impl	1.5-cdh5.5.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.5.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.5.2

CDH 5.5.1

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.5.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.5.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.5.1
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.5.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	260-mr1-cdh5.5.1
	org.apache.hadoop	hadoop-core	260-mr1-cdh5.5.1
	org.apache.hadoop	hadoop-examples	260-mr1-cdh5.5.1
	org.apache.hadoop	hadoop-minicluster	260-mr1-cdh5.5.1
	org.apache.hadoop	hadoop-streaming	260-mr1-cdh5.5.1
	org.apache.hadoop	hadoop-test	260-mr1-cdh5.5.1
	org.apache.hadoop	hadoop-tools	260-mr1-cdh5.5.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.5.1
	org.apache.hive	hive-ant	1.1.0-cdh5.5.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.5.1
	org.apache.hive	hive-cli	1.1.0-cdh5.5.1
	org.apache.hive	hive-common	1.1.0-cdh5.5.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.5.1
	org.apache.hive	hive-exec	1.1.0-cdh5.5.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.5.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.5.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.5.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.5.1
	org.apache.hive	hive-serde	1.1.0-cdh5.5.1
	org.apache.hive	hive-service	1.1.0-cdh5.5.1
	org.apache.hive	hive-shims	1.1.0-cdh5.5.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.5.1
	org.apache.hive	spark-client	1.1.0-cdh5.5.1
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.5.1
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.5.1

Project	groupId	artifactId	version
	org.apache.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.5.1
	org.apache.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.5.1
	org.apache.hcatalog	hive-webhcat	1.1.0-cdh5.5.1
	org.apache.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.5.1
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.5.1
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.5.1
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.5.1
HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-client	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-common	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-examples	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-it	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-rest	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-server	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-shell	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.5.1
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.5.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.5.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.5.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.5.1
	org.apache.pig	piggybank	0.12.0-cdh5.5.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.5.1
	org.apache.pig	pigunit	0.12.0-cdh5.5.1
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.5.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-core	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-node	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.5.1
	org.apache.flume	flume-tools	1.6.0-cdh5.5.1
	org.apache.flume	flume-file-channel	1.6.0-cdh5.5.1
	org.apache.flume	flume-jdbc-channel	1.6.0-cdh5.5.1
	org.apache.flume	flume-kafka-channel	1.6.0-cdh5.5.1
	org.apache.flume	flume-spillable-memory-channel	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-log4jappender	1.6.0-cdh5.5.1
	org.apache.flume	flume-avro-source	1.6.0-cdh5.5.1
	org.apache.flume	flume-thrift-source	1.6.0-cdh5.5.1
	org.apache.flume	flume-dataset-sink	1.6.0-cdh5.5.1
	org.apache.flume	flume-hdfs-sink	1.6.0-cdh5.5.1
	org.apache.flume	flume-hive-sink	1.6.0-cdh5.5.1
	org.apache.flume	flume-irc-sink	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-elasticsearch-sink	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-hbase-sink	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-kafka-sink	1.6.0-cdh5.5.1
	org.apache.flume	flume-ng-morphline-solr-sink	1.6.0-cdh5.5.1
	org.apache.flume	flume-jms-source	1.6.0-cdh5.5.1
	org.apache.flume	flume-kafka-source	1.6.0-cdh5.5.1
	org.apache.flume	flume-scribe-source	1.6.0-cdh5.5.1
	org.apache.flume	flume-taildir-source	1.6.0-cdh5.5.1
	org.apache.flume	flume-twitter-source	1.6.0-cdh5.5.1
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-core	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-examples	4.1.0-cdh5.5.1

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop	2.9.1-cdh5.5.1
	org.apache.oozie	oozie-hadoop-distcp	2.9.1-cdh5.5.1
	org.apache.oozie	oozie-hadoop-test	2.9.1-cdh5.5.1
	org.apache.oozie	oozie-hadoop-utils	2.9.1-cdh5.5.1
	org.apache.oozie	oozie-hbase	1.0-cdh5.5.1
	org.apache.oozie	oozie-hcatalog	1.0-cdh5.5.1
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.5.1
	org.apache.oozie	oozie-mini	4.1.0-cdh5.5.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.5.1
	org.apache.mahout	mahout-core	0.9-cdh5.5.1
	org.apache.mahout	mahout-examples	0.9-cdh5.5.1
	org.apache.mahout	mahout-integration	0.9-cdh5.5.1
	org.apache.mahout	mahout-math	0.9-cdh5.5.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.5.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.5.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.5.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.5.1
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.5.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.5.1
	org.apache.sqoop	test	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.5.1

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.5.1
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.5.1
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-dist	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.5.1
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.5.1
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.5.1
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.5.1
	com.twitter	parquet-cascading	1.5.0-cdh5.5.1
	com.twitter	parquet-column	1.5.0-cdh5.5.1

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-common	1.5.0-cdh5.5.1
	com.twitter	parquet-encoding	1.5.0-cdh5.5.1
	com.twitter	parquet-format	2.1.0-cdh5.5.1
	com.twitter	parquet-generator	1.5.0-cdh5.5.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.5.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.5.1
	com.twitter	parquet-jackson	1.5.0-cdh5.5.1
	com.twitter	parquet-pig	1.5.0-cdh5.5.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.5.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.5.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.5.1
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.5.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.5.1
	com.twitter	parquet-thrift	1.5.0-cdh5.5.1
	com.twitter	parquet-tools	1.5.0-cdh5.5.1
Llama	com.douderallama	llama	1.0.0-cdh5.5.1
Apache Spark	org.apache.spark	spark-assembly_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-bagel_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-catalyst_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-core_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-graphx_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-hive_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-launcher_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-mllib_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-network-common_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-network-shuffle_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-network-yarn_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-repl_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-sql_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.5.0-cdh5.5.1

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-flume_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-streaming-kafka_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-streaming-twitter_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-streaming_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-unsafe_2.10	1.5.0-cdh5.5.1
	org.apache.spark	spark-yarn_2.10	1.5.0-cdh5.5.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.5.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.5.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.5.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.5.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.5.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.5.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.5.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.5.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.5.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.5.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.5.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.5.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.5.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.5.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.5.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.5.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.5.1
	org.apache.avro	avro-tools	1.7.6-cdh5.5.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.5.1
	org.apache.avro	trevni-core	1.7.6-cdh5.5.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.5.1

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.5.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.5.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.5.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.5.1
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.5.1
	org.kitesdk	kite-data-s3	1.0.0-cdh5.5.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.5.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.5.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.5.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.5.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.5.1
	org.kitesdk	kite-tools	1.0.0-cdh5.5.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.5.1

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-core	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-join	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.5.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.5.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.5.1
	org.apache.solr	solr-cell	4.10.3-cdh5.5.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.5.1
	org.apache.solr	solr-core	4.10.3-cdh5.5.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.5.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.5.1
	org.apache.solr	solr-langid	4.10.3-cdh5.5.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.5.1
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.5.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.5.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.5.1
	org.apache.solr	solr-uima	4.10.3-cdh5.5.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.5.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.5.1
	com.cloudera.search	search-mr	1.0.0-cdh5.5.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.5.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.5.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.5.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.5.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.5.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.5.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.5.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.5.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.5.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.5.1
	com.ngdata	hbase-sep-api	1.5-cdh5.5.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.5.1
	com.ngdata	hbase-sep-impl	1.5-cdh5.5.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.5.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.5.1

CDH 5.5.0

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.5.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-archive-logs	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.5.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.5.0
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.5.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.5.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.5.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.5.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.5.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.5.0
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.5.0
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.5.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.5.0
	org.apache.hive	hive-ant	1.1.0-cdh5.5.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.5.0
	org.apache.hive	hive-cli	1.1.0-cdh5.5.0
	org.apache.hive	hive-common	1.1.0-cdh5.5.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.5.0
	org.apache.hive	hive-exec	1.1.0-cdh5.5.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.5.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.5.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.5.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.5.0
	org.apache.hive	hive-serde	1.1.0-cdh5.5.0

Project	groupId	artifactId	version
	org.apache.hive	hive-service	1.1.0-cdh5.5.0
	org.apache.hive	hive-shims	1.1.0-cdh5.5.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.5.0
	org.apache.hive	spark-client	1.1.0-cdh5.5.0
	org.apache.hcatalog	hive-hcatalog-core	1.1.0-cdh5.5.0
	org.apache.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.5.0
	org.apache.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.5.0
	org.apache.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.5.0
	org.apache.hcatalog	hive-webhcat	1.1.0-cdh5.5.0
	org.apache.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.5.0
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.5.0
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.5.0
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.5.0
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-client	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-common	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-examples	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-it	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-rest	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-server	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-shell	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.5.0
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.5.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.5.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.6-cdh5.5.0

Version and Download Information

Project	groupId	artifactId	version
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.5.0
	org.apache.pig	piggybank	0.12.0-cdh5.5.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.5.0
	org.apache.pig	pigunit	0.12.0-cdh5.5.0
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-configuration	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-core	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-embedded-agent	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-node	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-sdk	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-tests	1.6.0-cdh5.5.0
	org.apache.flume	flume-tools	1.6.0-cdh5.5.0
	org.apache.flume	flume-file-channel	1.6.0-cdh5.5.0
	org.apache.flume	flume-jdbc-channel	1.6.0-cdh5.5.0
	org.apache.flume	flume-kafka-channel	1.6.0-cdh5.5.0
	org.apache.flume	flume-spillable-memory-channel	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-log4jappender	1.6.0-cdh5.5.0
	org.apache.flume	flume-avro-source	1.6.0-cdh5.5.0
	org.apache.flume	flume-thrift-source	1.6.0-cdh5.5.0
	org.apache.flume	flume-dataset-sink	1.6.0-cdh5.5.0
	org.apache.flume	flume-hdfs-sink	1.6.0-cdh5.5.0
	org.apache.flume	flume-hive-sink	1.6.0-cdh5.5.0
	org.apache.flume	flume-irc-sink	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-elasticsearch-sink	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-hbase-sink	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-kafka-sink	1.6.0-cdh5.5.0
	org.apache.flume	flume-ng-morphline-solr-sink	1.6.0-cdh5.5.0
	org.apache.flume	flume-jms-source	1.6.0-cdh5.5.0
	org.apache.flume	flume-kafka-source	1.6.0-cdh5.5.0
	org.apache.flume	flume-scribe-source	1.6.0-cdh5.5.0

Project	groupId	artifactId	version
	org.apache.flume	flume-taildir-source	1.6.0-cdh5.5.0
	org.apache.flume	flume-twitter-source	1.6.0-cdh5.5.0
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-core	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-examples	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.5.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.5.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.5.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.5.0
	org.apache.mahout	mahout-core	0.9-cdh5.5.0
	org.apache.mahout	mahout-examples	0.9-cdh5.5.0
	org.apache.mahout	mahout-integration	0.9-cdh5.5.0
	org.apache.mahout	mahout-math	0.9-cdh5.5.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.5.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.5.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.5.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.5.0
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.5.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.5.0
	org.apache.sqoop	test	1.99.5-cdh5.5.0

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.5.0
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.5.0
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-binding-hive	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-binding-solr	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-core-common	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-core-model-db	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-core-model-indexer	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-core-model-search	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-dist	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-hdfs-common	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-hdfs-dist	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-hdfs-service	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-policy-common	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-policy-db	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-policy-indexer	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-policy-search	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-provider-cache	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-provider-common	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-provider-db	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-provider-file	1.5.1-cdh5.5.0
	org.apache.sentry	sentry-tests-hive	1.5.1-cdh5.5.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-tests-solr	1.5.1-cdh5.5.0
	org.apache.sentry	solr-sentry-handlers	1.5.1-cdh5.5.0
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.5.0
	com.twitter	parquet-cascading	1.5.0-cdh5.5.0
	com.twitter	parquet-column	1.5.0-cdh5.5.0
	com.twitter	parquet-common	1.5.0-cdh5.5.0
	com.twitter	parquet-encoding	1.5.0-cdh5.5.0
	com.twitter	parquet-format	2.1.0-cdh5.5.0
	com.twitter	parquet-generator	1.5.0-cdh5.5.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.5.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.5.0
	com.twitter	parquet-jackson	1.5.0-cdh5.5.0
	com.twitter	parquet-pig	1.5.0-cdh5.5.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.5.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.5.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.5.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.5.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.5.0
	com.twitter	parquet-thrift	1.5.0-cdh5.5.0
	com.twitter	parquet-tools	1.5.0-cdh5.5.0
Llama	com.doudera.llama	llama	1.0.0-cdh5.5.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-bagel_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-catalyst_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-core_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-graphx_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-hive_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-launcher_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-mllib_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-network-common_2.10	1.5.0-cdh5.5.0

Project	groupId	artifactId	version
	org.apache.spark	spark-network-shuffle_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-network-yarn_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-repl_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-sql_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-streaming-flume_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-streaming-kafka_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-streaming-mqtt-assembly_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-streaming-twitter_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-streaming_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-unsafe_2.10	1.5.0-cdh5.5.0
	org.apache.spark	spark-yarn_2.10	1.5.0-cdh5.5.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.5.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.5.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.5.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.5.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.5.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.5.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.5.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.5.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.5.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.5.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.5.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.5.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.5.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.5.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.5.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.5.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	avro-thrift	1.7.6-cdh5.5.0
	org.apache.avro	avro-tools	1.7.6-cdh5.5.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.5.0
	org.apache.avro	trevni-core	1.7.6-cdh5.5.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.5.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.5.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.5.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.5.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.5.0
	org.kitesdk	kite-data-oozie	1.0.0-cdh5.5.0
	org.kitesdk	kite-data-s3	1.0.0-cdh5.5.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.5.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.5.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.5.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.5.0

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.5.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.5.0
	org.kitesdk	kite-tools	1.0.0-cdh5.5.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-misc	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.5.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.5.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.5.0
	org.apache.solr	solr-cell	4.10.3-cdh5.5.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.5.0
	org.apache.solr	solr-core	4.10.3-cdh5.5.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.5.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.5.0
	org.apache.solr	solr-langid	4.10.3-cdh5.5.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.5.0
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.5.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.5.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.5.0
	org.apache.solr	solr-uima	4.10.3-cdh5.5.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.5.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.5.0
	com.cloudera.search	search-mr	1.0.0-cdh5.5.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.5.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.5.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.5.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.5.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.5.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.5.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.5.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.5.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.5.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.5.0
	com.ngdata	hbase-sep-api	1.5-cdh5.5.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.5.0
	com.ngdata	hbase-sep-impl	1.5-cdh5.5.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.5.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.5.0

Maven Artifacts for CDH 5.4.x Releases
CDH 5.4.11

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.11

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.11
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.11

Project	groupId	artifactId	version
	org.apache.hadoop.contrib	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.11
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.4.11
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.4.11
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.4.11
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.11
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.11
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.11
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.11
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.11
	org.apache.hive	hive-ant	1.1.0-cdh5.4.11
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.11
	org.apache.hive	hive-cli	1.1.0-cdh5.4.11
	org.apache.hive	hive-common	1.1.0-cdh5.4.11
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.11
	org.apache.hive	hive-exec	1.1.0-cdh5.4.11
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.11
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.11
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.11

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.11
	org.apache.hive	hive-serde	1.1.0-cdh5.4.11
	org.apache.hive	hive-service	1.1.0-cdh5.4.11
	org.apache.hive	hive-shims	1.1.0-cdh5.4.11
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.11
	org.apache.hive	spark-client	1.1.0-cdh5.4.11
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.4.11
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.11
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.4.11
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.4.11
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.4.11
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.4.11
	org.apache.hive.shims	hive-shims-0.23	1.1.0-cdh5.4.11
	org.apache.hive.shims	hive-shims-common	1.1.0-cdh5.4.11
	org.apache.hive.shims	hive-shims-scheduler	1.1.0-cdh5.4.11
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.11

Project	groupId	artifactId	version
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.11
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.11
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.11
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.11
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.11
	org.apache.pig	piggybank	0.12.0-cdh5.4.11
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.11
	org.apache.pig	pigunit	0.12.0-cdh5.4.11
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.11
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.11
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.11
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.11
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.11
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.11

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.11
	org.apache.flume	flume-tools	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-channels	flume-file-channel	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-channels	flume-jdbc-channel	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-channels	flume-kafka-channel	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-channels	flume-spillable-memory-channel	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-clients	flume-ng-log4jappender	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sources	flume-avro-source	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sources	flume-thrift-source	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sinks	flume-dataset-sink	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sinks	flume-hdfs-sink	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sinks	flume-hive-sink	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sinks	flume-irc-sink	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sinks	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sinks	flume-ng-hbase-sink	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sinks	flume-ng-kafka-sink	1.5.0-cdh5.4.11

Project	groupId	artifactId	version
	org.apache.flume.flume-ng-sinks	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sources	flume-jms-source	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sources	flume-kafka-source	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sources	flume-scribe-source	1.5.0-cdh5.4.11
	org.apache.flume.flume-ng-sources	flume-twitter-source	1.5.0-cdh5.4.11
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-hadoop	2.6.0-mr1-cdh5.4.11.oozie-4.1.0-cdh5.4.11
	org.apache.oozie	oozie-hadoop-distcp	2.6.0-mr1-cdh5.4.11.oozie-4.1.0-cdh5.4.11
	org.apache.oozie	oozie-hadoop-test	2.6.0-mr1-cdh5.4.11.oozie-4.1.0-cdh5.4.11
	org.apache.oozie	oozie-hadoop-utils	2.6.0-mr1-cdh5.4.11.oozie-4.1.0-cdh5.4.11
	org.apache.oozie	oozie-hbase	1.0.0-cdh5.4.11.oozie-4.1.0-cdh5.4.11
	org.apache.oozie	oozie-hcatalog	1.1.0-cdh5.4.11.oozie-4.1.0-cdh5.4.11
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.11

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.11
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.11
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.4.11
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.11
	org.apache.mahout	mahout-core	0.9-cdh5.4.11
	org.apache.mahout	mahout-examples	0.9-cdh5.4.11
	org.apache.mahout	mahout-integration	0.9-cdh5.4.11
	org.apache.mahout	mahout-math	0.9-cdh5.4.11
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.11
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.11

Project	groupId	artifactId	version
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.11
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.11
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.4.11
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.11
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.11
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.11
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.11
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.11

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.11
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.11
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.11
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.11
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.11
	org.apache.sqoop	test	1.99.5-cdh5.4.11
	org.apache.sqoop.connector	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.11
	org.apache.sqoop.connector	sqoop-connector-hdfs	1.99.5-cdh5.4.11
	org.apache.sqoop.connector	sqoop-connector-kafka	1.99.5-cdh5.4.11
	org.apache.sqoop.connector	sqoop-connector-kite	1.99.5-cdh5.4.11
	org.apache.sqoop.execution	sqoop-execution-mapreduce	1.99.5-cdh5.4.11
	org.apache.sqoop.repository	sqoop-repository-common	1.99.5-cdh5.4.11
	org.apache.sqoop.repository	sqoop-repository-derby	1.99.5-cdh5.4.11
	org.apache.sqoop.repository	sqoop-repository-postgresql	1.99.5-cdh5.4.11
	org.apache.sqoop.submission	sqoop-submission-mapreduce	1.99.5-cdh5.4.11
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.11

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.11
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.11
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.11
	com.twitter	parquet-cascading	1.5.0-cdh5.4.11
	com.twitter	parquet-column	1.5.0-cdh5.4.11

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-common	1.5.0-cdh5.4.11
	com.twitter	parquet-encoding	1.5.0-cdh5.4.11
	com.twitter	parquet-format	2.1.0-cdh5.4.11
	com.twitter	parquet-generator	1.5.0-cdh5.4.11
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.11
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.11
	com.twitter	parquet-jackson	1.5.0-cdh5.4.11
	com.twitter	parquet-pig	1.5.0-cdh5.4.11
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.11
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.11
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.11
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.11
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.11
	com.twitter	parquet-thrift	1.5.0-cdh5.4.11
	com.twitter	parquet-tools	1.5.0-cdh5.4.11
Llama	com.cloudera.llama	llama	1.0.0-cdh5.4.11
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.11

Project	groupId	artifactId	version
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.11
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.11
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.11
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.11
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.11
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.11
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.11

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.11
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.11
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.11
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.11
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.11
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.11
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.11
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.11
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.11
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.11
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.11
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.11
	org.apache.avro	avro-tools	1.7.6-cdh5.4.11
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.11
	org.apache.avro	trevni-core	1.7.6-cdh5.4.11
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.11
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.11
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.11
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.11
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.11
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.11

Project	groupId	artifactId	version
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.11
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.11
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.11
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.11
	org.kitesdk	kite-tools	1.0.0-cdh5.4.11

Version and Download Information

Project	groupId	artifactId	version
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.11

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.11
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.11
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.11
	org.apache.solr	solr-cell	4.10.3-cdh5.4.11
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.11
	org.apache.solr	solr-core	4.10.3-cdh5.4.11
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.11
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.11
	org.apache.solr	solr-langid	4.10.3-cdh5.4.11
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.11
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.11
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.11
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.11
	org.apache.solr	solr-uima	4.10.3-cdh5.4.11
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.11
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.11
	com.cloudera.search	search-mr	1.0.0-cdh5.4.11
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.11

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.11
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.11
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.11
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.11
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.11
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.11
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.11
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.11
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.11
	com.ngdata	hbase-sep-api	1.5-cdh5.4.11
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.11
	com.ngdata	hbase-sep-impl	1.5-hbase1.0-cdh5.4.11
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.11
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.11

CDH 5.4.10

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.10

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.10
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.10
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	260mr1-cdh5.4.10
	org.apache.hadoop	hadoop-core	260mr1-cdh5.4.10
	org.apache.hadoop	hadoop-examples	260mr1-cdh5.4.10

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.10
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.10
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.10
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.10
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.10
	org.apache.hive	hive-ant	1.1.0-cdh5.4.10
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.10
	org.apache.hive	hive-cli	1.1.0-cdh5.4.10
	org.apache.hive	hive-common	1.1.0-cdh5.4.10
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.10
	org.apache.hive	hive-exec	1.1.0-cdh5.4.10
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.10
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.10
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.10
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.10
	org.apache.hive	hive-serde	1.1.0-cdh5.4.10
	org.apache.hive	hive-service	1.1.0-cdh5.4.10
	org.apache.hive	hive-shims	1.1.0-cdh5.4.10
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.10
	org.apache.hive	spark-client	1.1.0-cdh5.4.10
	org.apache.hive.hcatalog	hive-hcatalog-core	1.1.0-cdh5.4.10
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.10
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.4.10
	org.apache.hive.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.4.10
	org.apache.hive.hcatalog	hive-webhcat	1.1.0-cdh5.4.10
	org.apache.hive.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.4.10
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.4.10
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.4.10
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.4.10
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.10

Project	groupId	artifactId	version
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.10
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.10
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.10
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.10
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.10
	org.apache.pig	piggybank	0.12.0-cdh5.4.10
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.10
	org.apache.pig	pigunit	0.12.0-cdh5.4.10
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.10
	org.apache.flume	flume-tools	1.5.0-cdh5.4.10
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.10

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.10
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.10
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.10
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.10
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.10
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.10
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.10
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.10
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.10
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.10
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.10
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.10
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.10
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.10
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.10

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.10
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.10
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.4.10
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.10
	org.apache.mahout	mahout-core	0.9-cdh5.4.10
	org.apache.mahout	mahout-examples	0.9-cdh5.4.10
	org.apache.mahout	mahout-integration	0.9-cdh5.4.10
	org.apache.mahout	mahout-math	0.9-cdh5.4.10
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.10
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.10

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.10
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.10
Apache DataFu (Incubating)	com.linkedin.datafu	datafu	1.1.0-cdh5.4.10
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.10
	org.apache.sqoop	test	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.4.10
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.4.10
Apache Sentry (Incubating)	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.10

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.10
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.10
Apache Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.10
	com.twitter	parquet-cascading	1.5.0-cdh5.4.10
	com.twitter	parquet-column	1.5.0-cdh5.4.10
	com.twitter	parquet-common	1.5.0-cdh5.4.10
	com.twitter	parquet-encoding	1.5.0-cdh5.4.10
	com.twitter	parquet-format	2.1.0-cdh5.4.10
	com.twitter	parquet-generator	1.5.0-cdh5.4.10
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.10
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.10
	com.twitter	parquet-jackson	1.5.0-cdh5.4.10
	com.twitter	parquet-pig	1.5.0-cdh5.4.10
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.10
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.10

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.10
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.10
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.10
	com.twitter	parquet-thrift	1.5.0-cdh5.4.10
	com.twitter	parquet-tools	1.5.0-cdh5.4.10
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.10
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.10
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.10
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.10
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.10
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.10

Project	groupId	artifactId	version
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.10
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.10
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.10
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.10
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.10
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.10
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.10
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.10
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.10
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.10
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.10
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.10
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.10
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.10
	org.apache.avro	avro-tools	1.7.6-cdh5.4.10
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.10
	org.apache.avro	trevni-core	1.7.6-cdh5.4.10
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.10
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.10
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.10
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.10
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.10
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.10
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.10
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.10
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.10

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.10
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.10
	org.kitesdk	kite-tools	1.0.0-cdh5.4.10
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.10

Project	groupId	artifactId	version
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.10
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.10
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.10
	org.apache.solr	solr-cell	4.10.3-cdh5.4.10
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.10
	org.apache.solr	solr-core	4.10.3-cdh5.4.10
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.10
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.10
	org.apache.solr	solr-langid	4.10.3-cdh5.4.10
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.10
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.10
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.10
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.10
	org.apache.solr	solr-uima	4.10.3-cdh5.4.10
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.10
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.10
	com.cloudera.search	search-mr	1.0.0-cdh5.4.10
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.10
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.10
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.10
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.10

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.10
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.10
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.10
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.10
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.10
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.10
	com.ngdata	hbase-sep-api	1.5-cdh5.4.10
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.10
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.10
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.10
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.10

CDH 5.4.9

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.9

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.9
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.9
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.4.9
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.4.9
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.4.9
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.9

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.9
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.9
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.9
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.9
	org.apache.hive	hive-ant	1.1.0-cdh5.4.9
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.9
	org.apache.hive	hive-cli	1.1.0-cdh5.4.9
	org.apache.hive	hive-common	1.1.0-cdh5.4.9
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.9
	org.apache.hive	hive-exec	1.1.0-cdh5.4.9
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.9
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.9
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.9
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.9
	org.apache.hive	hive-serde	1.1.0-cdh5.4.9
	org.apache.hive	hive-service	1.1.0-cdh5.4.9
	org.apache.hive	hive-shims	1.1.0-cdh5.4.9
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.9
	org.apache.hive	spark-client	1.1.0-cdh5.4.9
	org.apache.hive	hive-hcatalog	1.1.0-cdh5.4.9
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.4.9
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.9
	org.apache.hive	hive-hcatalog-server-extensions	1.1.0-cdh5.4.9
	org.apache.hive	hive-hcatalog-streaming	1.1.0-cdh5.4.9
	org.apache.hive	hive-webhcat	1.1.0-cdh5.4.9
	org.apache.hive	hive-webhcat-java-client	1.1.0-cdh5.4.9
	org.apache.hive	hive-shims-0.23	1.1.0-cdh5.4.9
	org.apache.hive	hive-shims-common	1.1.0-cdh5.4.9
	org.apache.hive	hive-shims-scheduler	1.1.0-cdh5.4.9
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.9

Project	groupId	artifactId	version
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.9
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.9
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.9
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.9
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.9
	org.apache.pig	piggybank	0.12.0-cdh5.4.9
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.9
	org.apache.pig	pigunit	0.12.0-cdh5.4.9
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.9
	org.apache.flume	flume-tools	1.5.0-cdh5.4.9
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.9
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.9

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.9
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.9
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.9
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.9
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.9
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.9
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.9
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.9
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.9
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.9
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.9
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.9
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.9
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.9

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.9
	org.apache.oozie	oozie-mini	4.1.0-cdh5.4.9
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.9
	org.apache.mahout	mahout-core	0.9-cdh5.4.9
	org.apache.mahout	mahout-examples	0.9-cdh5.4.9
	org.apache.mahout	mahout-integration	0.9-cdh5.4.9
	org.apache.mahout	mahout-math	0.9-cdh5.4.9
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.9
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.9

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.9
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.9
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.4.9
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.9
	org.apache.sqoop	test	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.4.9
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.4.9
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.9

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.9
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.9
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.9
	com.twitter	parquet-cascading	1.5.0-cdh5.4.9
	com.twitter	parquet-column	1.5.0-cdh5.4.9
	com.twitter	parquet-common	1.5.0-cdh5.4.9
	com.twitter	parquet-encoding	1.5.0-cdh5.4.9
	com.twitter	parquet-format	2.1.0-cdh5.4.9
	com.twitter	parquet-generator	1.5.0-cdh5.4.9
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.9
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.9
	com.twitter	parquet-jackson	1.5.0-cdh5.4.9
	com.twitter	parquet-pig	1.5.0-cdh5.4.9
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.9
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.9
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.9
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.9

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.9
	com.twitter	parquet-thrift	1.5.0-cdh5.4.9
	com.twitter	parquet-tools	1.5.0-cdh5.4.9
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.9
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.9
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.9
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.9
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.9
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.9
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.9
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.9

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.9
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.9
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.9
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.9
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.9
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.9
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.9
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.9
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.9
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.9
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.9
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.9
	org.apache.avro	avro-tools	1.7.6-cdh5.4.9
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.9
	org.apache.avro	trevni-core	1.7.6-cdh5.4.9
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.9
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.9
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.9
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.9
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.9
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.9
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.9
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.9
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.9

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.9
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.9
	org.kitesdk	kite-tools	1.0.0-cdh5.4.9
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.9

Project	groupId	artifactId	version
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.9
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.9
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.9
	org.apache.solr	solr-cell	4.10.3-cdh5.4.9
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.9
	org.apache.solr	solr-core	4.10.3-cdh5.4.9
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.9
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.9
	org.apache.solr	solr-langid	4.10.3-cdh5.4.9
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.9
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.9
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.9
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.9
	org.apache.solr	solr-uima	4.10.3-cdh5.4.9
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.9
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.9
	com.cloudera.search	search-mr	1.0.0-cdh5.4.9
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.9
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.9
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.9
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.9
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.9
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.9

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.9
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.9
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.9
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.9
	com.ngdata	hbase-sep-api	1.5-cdh5.4.9
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.9
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.9
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.9
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.9

CDH 5.4.8

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.8

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.8
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.8
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.4.8
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.4.8
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.4.8
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.8
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.8
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.8

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.8
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.8
	org.apache.hive	hive-ant	1.1.0-cdh5.4.8
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.8
	org.apache.hive	hive-cli	1.1.0-cdh5.4.8
	org.apache.hive	hive-common	1.1.0-cdh5.4.8
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.8
	org.apache.hive	hive-exec	1.1.0-cdh5.4.8
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.8
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.8
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.8
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.8
	org.apache.hive	hive-serde	1.1.0-cdh5.4.8
	org.apache.hive	hive-service	1.1.0-cdh5.4.8
	org.apache.hive	hive-shims	1.1.0-cdh5.4.8
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.8
	org.apache.hive	spark-client	1.1.0-cdh5.4.8
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.4.8
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.8
	org.apache.hive	hive-hcatalog-server-extensions	1.1.0-cdh5.4.8
	org.apache.hive	hive-hcatalog-streaming	1.1.0-cdh5.4.8
	org.apache.hive	hive-webhcat	1.1.0-cdh5.4.8
	org.apache.hive	hive-webhcat-java-client	1.1.0-cdh5.4.8
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.4.8
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.4.8
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.4.8
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.8

Project	groupId	artifactId	version
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.8
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.8
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.8
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.8
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.8
	org.apache.pig	piggybank	0.12.0-cdh5.4.8
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.8
	org.apache.pig	pigunit	0.12.0-cdh5.4.8
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.8
	org.apache.flume	flume-tools	1.5.0-cdh5.4.8
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.8
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.8
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.8
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.8

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.8
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.8
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.8
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.8
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.8
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.8
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.8
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.8
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.8
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.8
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.8
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.8
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.8

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.8
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.8
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.4.8
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.8
	org.apache.mahout	mahout-core	0.9-cdh5.4.8
	org.apache.mahout	mahout-examples	0.9-cdh5.4.8
	org.apache.mahout	mahout-integration	0.9-cdh5.4.8
	org.apache.mahout	mahout-math	0.9-cdh5.4.8
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.8
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.8
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.8

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.8
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.4.8
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.8
	org.apache.sqoop	test	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.4.8
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.4.8
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.8

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.8
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.8
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.8
	com.twitter	parquet-cascading	1.5.0-cdh5.4.8
	com.twitter	parquet-column	1.5.0-cdh5.4.8
	com.twitter	parquet-common	1.5.0-cdh5.4.8
	com.twitter	parquet-encoding	1.5.0-cdh5.4.8
	com.twitter	parquet-format	2.1.0-cdh5.4.8
	com.twitter	parquet-generator	1.5.0-cdh5.4.8
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.8
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.8
	com.twitter	parquet-jackson	1.5.0-cdh5.4.8
	com.twitter	parquet-pig	1.5.0-cdh5.4.8
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.8
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.8
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.8
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.8
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.8
	com.twitter	parquet-thrift	1.5.0-cdh5.4.8

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-tools	1.5.0-cdh5.4.8
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.8
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.8
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.8
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.8
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.8
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.8
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.8
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.8
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.8
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.8

Project	groupId	artifactId	version
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.8
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.8
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.8
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.8
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.8
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.8
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.8
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.8
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.8
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.8
	org.apache.avro	avro-tools	1.7.6-cdh5.4.8
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.8
	org.apache.avro	trevni-core	1.7.6-cdh5.4.8
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.8
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.8
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.8
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.8
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.8
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.8
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.8
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.8
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.8

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.8
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.8
	org.kitesdk	kite-tools	1.0.0-cdh5.4.8
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.8

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.8
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.8
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.8
	org.apache.solr	solr-cell	4.10.3-cdh5.4.8
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.8
	org.apache.solr	solr-core	4.10.3-cdh5.4.8
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.8
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.8
	org.apache.solr	solr-langid	4.10.3-cdh5.4.8
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.8
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.8
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.8
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.8
	org.apache.solr	solr-uima	4.10.3-cdh5.4.8
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.8
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.8
	com.cloudera.search	search-mr	1.0.0-cdh5.4.8
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.8
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.8
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.8
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.8
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.8
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.8
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.8
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.8

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.8
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.8
	com.ngdata	hbase-sep-api	1.5-cdh5.4.8
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.8
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.8
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.8
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.8

CDH 5.4.7

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.7

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.7
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.7
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.4.7
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.4.7
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.4.7
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.7
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.7
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.7
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.7
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.7

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-ant	1.1.0-cdh5.4.7
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.7
	org.apache.hive	hive-cli	1.1.0-cdh5.4.7
	org.apache.hive	hive-common	1.1.0-cdh5.4.7
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.7
	org.apache.hive	hive-exec	1.1.0-cdh5.4.7
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.7
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.7
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.7
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.7
	org.apache.hive	hive-serde	1.1.0-cdh5.4.7
	org.apache.hive	hive-service	1.1.0-cdh5.4.7
	org.apache.hive	hive-shims	1.1.0-cdh5.4.7
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.7
	org.apache.hive	spark-client	1.1.0-cdh5.4.7
	org.apache.hcatalog	hive-hcatalog-core	1.1.0-cdh5.4.7
	org.apache.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.7
	org.apache.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.4.7
	org.apache.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.4.7
	org.apache.hcatalog	hive-webhcat	1.1.0-cdh5.4.7
	org.apache.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.4.7
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.4.7
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.4.7
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.4.7
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.7

Project	groupId	artifactId	version
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.7
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.7
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.7
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.7
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.7
	org.apache.pig	piggybank	0.12.0-cdh5.4.7
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.7
	org.apache.pig	pigunit	0.12.0-cdh5.4.7
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.7
	org.apache.flume	flume-tools	1.5.0-cdh5.4.7
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.7
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.7
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.7
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.7
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.7

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.7
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.7
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.7
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.7
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.7
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.7
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.7
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.7
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.7
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.7
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.7

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.7
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.7
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.4.7
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.7
	org.apache.mahout	mahout-core	0.9-cdh5.4.7
	org.apache.mahout	mahout-examples	0.9-cdh5.4.7
	org.apache.mahout	mahout-integration	0.9-cdh5.4.7
	org.apache.mahout	mahout-math	0.9-cdh5.4.7
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.7
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.7
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.7
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.4.7

Version and Download Information

Project	groupId	artifactId	version
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.7
	org.apache.sqoop	test	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.4.7
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.4.7
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.7

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.7
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.7
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.7
	com.twitter	parquet-cascading	1.5.0-cdh5.4.7
	com.twitter	parquet-column	1.5.0-cdh5.4.7
	com.twitter	parquet-common	1.5.0-cdh5.4.7
	com.twitter	parquet-encoding	1.5.0-cdh5.4.7
	com.twitter	parquet-format	2.1.0-cdh5.4.7
	com.twitter	parquet-generator	1.5.0-cdh5.4.7
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.7
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.7
	com.twitter	parquet-jackson	1.5.0-cdh5.4.7
	com.twitter	parquet-pig	1.5.0-cdh5.4.7
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.7
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.7
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.7
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.7
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.7
	com.twitter	parquet-thrift	1.5.0-cdh5.4.7
	com.twitter	parquet-tools	1.5.0-cdh5.4.7
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.7

Version and Download Information

Project	groupId	artifactId	version
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.7
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.7
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.7
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.7
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.7
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.7
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.7
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.7
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.7
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.7
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.7

Project	groupId	artifactId	version
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.7
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.7
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.7
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.7
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.7
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.7
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.7
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.7
	org.apache.avro	avro-tools	1.7.6-cdh5.4.7
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.7
	org.apache.avro	trevni-core	1.7.6-cdh5.4.7
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.7
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.7
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.7
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.7
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.7
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.7
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.7
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.7
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.7

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.7
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.7
	org.kitesdk	kite-tools	1.0.0-cdh5.4.7
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.7

Project	groupId	artifactId	version
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.7
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.7
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.7
	org.apache.solr	solr-cell	4.10.3-cdh5.4.7
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.7
	org.apache.solr	solr-core	4.10.3-cdh5.4.7
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.7
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.7
	org.apache.solr	solr-langid	4.10.3-cdh5.4.7
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.7
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.7
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.7
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.7
	org.apache.solr	solr-uima	4.10.3-cdh5.4.7
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.7
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.7
	com.cloudera.search	search-mr	1.0.0-cdh5.4.7
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.7
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.7
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.7
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.7
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.7
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.7
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.7
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.7
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.7
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.7

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-api	1.5-cdh5.4.7
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.7
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.7
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.7
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.7

CDH 5.4.5

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.5

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.5
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.5
Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.4.5
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.4.5
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.4.5
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.5
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.5
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.5
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.5
Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.5
	org.apache.hive	hive-ant	1.1.0-cdh5.4.5
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.5
	org.apache.hive	hive-cli	1.1.0-cdh5.4.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-common	1.1.0-cdh5.4.5
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.5
	org.apache.hive	hive-exec	1.1.0-cdh5.4.5
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.5
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.5
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.5
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.5
	org.apache.hive	hive-serde	1.1.0-cdh5.4.5
	org.apache.hive	hive-service	1.1.0-cdh5.4.5
	org.apache.hive	hive-shims	1.1.0-cdh5.4.5
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.5
	org.apache.hive	spark-client	1.1.0-cdh5.4.5
	org.apache.hcatalog	hive-hcatalog-core	1.1.0-cdh5.4.5
	org.apache.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.5
	org.apache.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.4.5
	org.apache.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.4.5
	org.apache.hcatalog	hive-webhcat	1.1.0-cdh5.4.5
	org.apache.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.4.5
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.4.5
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.4.5
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.4.5
HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.5

Project	groupId	artifactId	version
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.5
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.5
ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.5
Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.5
Pig	org.apache.pig	pig	0.12.0-cdh5.4.5
	org.apache.pig	piggybank	0.12.0-cdh5.4.5
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.5
	org.apache.pig	pigunit	0.12.0-cdh5.4.5
Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.5
	org.apache.flume	flume-tools	1.5.0-cdh5.4.5
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.5
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.5
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.5
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.5
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.5
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.5
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.5
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.5
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.5
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.5
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.5
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.5
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.5
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.5
Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.5
	org.apache.oozie	oozie-mini	4.1.0-cdh5.4.5
Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.5

Project	groupId	artifactId	version
	org.apache.mahout	mahout-core	0.9-cdh5.4.5
	org.apache.mahout	mahout-examples	0.9-cdh5.4.5
	org.apache.mahout	mahout-integration	0.9-cdh5.4.5
	org.apache.mahout	mahout-math	0.9-cdh5.4.5
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.5
Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.5
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.5
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.4.5
Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.5
	org.apache.sqoop	test	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.4.5
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.4.5
Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.5

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.5
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.5
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.5
	com.twitter	parquet-cascading	1.5.0-cdh5.4.5
	com.twitter	parquet-column	1.5.0-cdh5.4.5
	com.twitter	parquet-common	1.5.0-cdh5.4.5
	com.twitter	parquet-encoding	1.5.0-cdh5.4.5
	com.twitter	parquet-format	2.1.0-cdh5.4.5
	com.twitter	parquet-generator	1.5.0-cdh5.4.5
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.5
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.5
	com.twitter	parquet-jackson	1.5.0-cdh5.4.5
	com.twitter	parquet-pig	1.5.0-cdh5.4.5
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.5
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.5
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.5
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.5
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.5
	com.twitter	parquet-thrift	1.5.0-cdh5.4.5
	com.twitter	parquet-tools	1.5.0-cdh5.4.5
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.5
Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.5
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.5
Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.5
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.5
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.5
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.5
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.5
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.5
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.5
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.5
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.5
Avro	org.apache.avro	avro	1.7.6-cdh5.4.5
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.5
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.5
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.5

Project	groupId	artifactId	version
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.5
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.5
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.5
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.5
	org.apache.avro	avro-tools	1.7.6-cdh5.4.5
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.5
	org.apache.avro	trevni-core	1.7.6-cdh5.4.5
Kite	org.kitesdk	kite-data-core	1.0.0-cdh5.4.5
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.5
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.5
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.5
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.5
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.5
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.5
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.5
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.5

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.5
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.5
	org.kitesdk	kite-tools	1.0.0-cdh5.4.5
Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.5
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.5

Project	groupId	artifactId	version
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.5
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.5
	org.apache.solr	solr-cell	4.10.3-cdh5.4.5
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.5
	org.apache.solr	solr-core	4.10.3-cdh5.4.5
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.5
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.5
	org.apache.solr	solr-langid	4.10.3-cdh5.4.5
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.5
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.5
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.5
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.5
	org.apache.solr	solr-uima	4.10.3-cdh5.4.5
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.5
Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.5
	com.cloudera.search	search-mr	1.0.0-cdh5.4.5
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.5
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.5
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.5
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.5
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.5
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.5
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.5
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.5
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.5
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.5
	com.ngdata	hbase-sep-api	1.5-cdh5.4.5
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.5
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.5
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.5

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.5

CDH 5.4.4

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.4

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.4
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.4
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.4.4
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.4.4
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.4.4
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.4
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.4
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.4
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.4
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.4
	org.apache.hive	hive-ant	1.1.0-cdh5.4.4
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.4
	org.apache.hive	hive-cli	1.1.0-cdh5.4.4
	org.apache.hive	hive-common	1.1.0-cdh5.4.4
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.4
	org.apache.hive	hive-exec	1.1.0-cdh5.4.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.4
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.4
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.4
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.4
	org.apache.hive	hive-serde	1.1.0-cdh5.4.4
	org.apache.hive	hive-service	1.1.0-cdh5.4.4
	org.apache.hive	hive-shims	1.1.0-cdh5.4.4
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.4
	org.apache.hive	spark-client	1.1.0-cdh5.4.4
	org.apache.hive	hive-hcatalog	1.1.0-cdh5.4.4
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.4.4
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.4
	org.apache.hive	hive-hcatalog-server-extensions	1.1.0-cdh5.4.4
	org.apache.hive	hive-hcatalog-streaming	1.1.0-cdh5.4.4
	org.apache.hive	hive-webhcat	1.1.0-cdh5.4.4
	org.apache.hive	hive-webhcat-java-client	1.1.0-cdh5.4.4
	org.apache.hive	hive-shims	1.1.0-cdh5.4.4
	org.apache.hive	hive-shims-0.23	1.1.0-cdh5.4.4
	org.apache.hive	hive-shims-common	1.1.0-cdh5.4.4
	org.apache.hive	hive-shims-scheduler	1.1.0-cdh5.4.4
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.4

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.4
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.4
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.4
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.4
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.4
	org.apache.pig	piggybank	0.12.0-cdh5.4.4
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.4
	org.apache.pig	pigunit	0.12.0-cdh5.4.4
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.4
	org.apache.flume	flume-tools	1.5.0-cdh5.4.4
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.4
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.4
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.4
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.4
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.4
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.4
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.4
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.4
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.4
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.4
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.4
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.4
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.4
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.4
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.4
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.4
	org.apache.oozie	oozie-mini	4.1.0-cdh5.4.4
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.4
	org.apache.mahout	mahout-core	0.9-cdh5.4.4
	org.apache.mahout	mahout-examples	0.9-cdh5.4.4

Project	groupId	artifactId	version
	org.apache.mahout	mahout-integration	0.9-cdh5.4.4
	org.apache.mahout	mahout-math	0.9-cdh5.4.4
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.4
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.4
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.4
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.4.4
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.4
	org.apache.sqoop	test	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.4.4
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.4.4
Apache Sentry	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.4

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.4
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.4
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.4
	com.twitter	parquet-cascading	1.5.0-cdh5.4.4
	com.twitter	parquet-column	1.5.0-cdh5.4.4
	com.twitter	parquet-common	1.5.0-cdh5.4.4
	com.twitter	parquet-encoding	1.5.0-cdh5.4.4
	com.twitter	parquet-format	2.1.0-cdh5.4.4
	com.twitter	parquet-generator	1.5.0-cdh5.4.4
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.4
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.4
	com.twitter	parquet-jackson	1.5.0-cdh5.4.4
	com.twitter	parquet-pig	1.5.0-cdh5.4.4
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.4
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.4
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.4
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.4
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.4
	com.twitter	parquet-thrift	1.5.0-cdh5.4.4
	com.twitter	parquet-tools	1.5.0-cdh5.4.4
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.4
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.4
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.4
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.4
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.4
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.4
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.4
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.4
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.4
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.4
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.4
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.4
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.4
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.4
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.4
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.4
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.4
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.4

Project	groupId	artifactId	version
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.4
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.4
	org.apache.avro	avro-tools	1.7.6-cdh5.4.4
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.4
	org.apache.avro	trevni-core	1.7.6-cdh5.4.4
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.4
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.4
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.4
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.4
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.4
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.4
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.4
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.4
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.4
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.4

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.4
	org.kitesdk	kite-tools	1.0.0-cdh5.4.4
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.4
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.4
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.4

Project	groupId	artifactId	version
	org.apache.solr	solr-cell	4.10.3-cdh5.4.4
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.4
	org.apache.solr	solr-core	4.10.3-cdh5.4.4
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.4
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.4
	org.apache.solr	solr-langid	4.10.3-cdh5.4.4
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.4
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.4
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.4
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.4
	org.apache.solr	solr-uima	4.10.3-cdh5.4.4
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.4
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.4
	com.cloudera.search	search-mr	1.0.0-cdh5.4.4
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.4
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.4
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.4
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.4
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.4
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.4
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.4
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.4
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.4
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.4
	com.ngdata	hbase-sep-api	1.5-cdh5.4.4
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.4
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.4
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.4
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.4

Version and Download Information

CDH 5.4.3

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.3

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.3
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-core	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-examples	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-minicluster	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-streaming	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-test	2.6.0-cdh5.4.3
	org.apache.hadoop	hadoop-tools	2.6.0-cdh5.4.3
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.3
	org.apache.hive	hive-ant	1.1.0-cdh5.4.3
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.3
	org.apache.hive	hive-cli	1.1.0-cdh5.4.3
	org.apache.hive	hive-common	1.1.0-cdh5.4.3
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.3
	org.apache.hive	hive-exec	1.1.0-cdh5.4.3
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.3
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.3
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.3
	org.apache.hive	hive-serde	1.1.0-cdh5.4.3
	org.apache.hive	hive-service	1.1.0-cdh5.4.3
	org.apache.hive	hive-shims	1.1.0-cdh5.4.3
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.3
	org.apache.hive	spark-client	1.1.0-cdh5.4.3
	org.apache.hive	hive-hcatalog	1.1.0-cdh5.4.3
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.4.3
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.3
	org.apache.hive	hive-hcatalog-server-extensions	1.1.0-cdh5.4.3
	org.apache.hive	hive-hcatalog-streaming	1.1.0-cdh5.4.3
	org.apache.hive	hive-webhcat	1.1.0-cdh5.4.3
	org.apache.hive	hive-webhcat-java-client	1.1.0-cdh5.4.3
	org.apache.hive	hive-shims-0.23	1.1.0-cdh5.4.3
	org.apache.hive	hive-shims-common	1.1.0-cdh5.4.3
	org.apache.hive	hive-shims-scheduler	1.1.0-cdh5.4.3
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.3
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.3

Project	groupId	artifactId	version
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.3
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.3
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.3
	org.apache.pig	piggybank	0.12.0-cdh5.4.3
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.3
	org.apache.pig	pigunit	0.12.0-cdh5.4.3
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.3
	org.apache.flume	flume-tools	1.5.0-cdh5.4.3
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.3
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.3
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.3
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.3
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.3
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.3
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.3
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.3
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.3
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.3
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.3
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.3
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.3
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.3
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.3
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.3
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.4.3
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.3
	org.apache.mahout	mahout-core	0.9-cdh5.4.3
	org.apache.mahout	mahout-examples	0.9-cdh5.4.3
	org.apache.mahout	mahout-integration	0.9-cdh5.4.3
	org.apache.mahout	mahout-math	0.9-cdh5.4.3
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.3

Project	groupId	artifactId	version
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.3
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.3
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.4.3
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.3
	org.apache.sqoop	test	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.4.3
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.4.3
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.3

Project	groupId	artifactId	version
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.3
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.3
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.3
	com.twitter	parquet-cascading	1.5.0-cdh5.4.3
	com.twitter	parquet-column	1.5.0-cdh5.4.3
	com.twitter	parquet-common	1.5.0-cdh5.4.3
	com.twitter	parquet-encoding	1.5.0-cdh5.4.3
	com.twitter	parquet-format	2.1.0-cdh5.4.3
	com.twitter	parquet-generator	1.5.0-cdh5.4.3
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.3
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.3
	com.twitter	parquet-jackson	1.5.0-cdh5.4.3
	com.twitter	parquet-pig	1.5.0-cdh5.4.3
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.3
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.3
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.3
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.3
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.3
	com.twitter	parquet-thrift	1.5.0-cdh5.4.3
	com.twitter	parquet-tools	1.5.0-cdh5.4.3
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.3
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.3
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.3
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.3
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.3
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.3
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.3
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.3
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.3
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.3
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.3
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.3
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.3
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.3
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.3
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.3
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.3
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.3
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.3
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.3
	org.apache.avro	avro-tools	1.7.6-cdh5.4.3

Project	groupId	artifactId	version
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.3
	org.apache.avro	trevni-core	1.7.6-cdh5.4.3
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.3
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.3
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.3
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.3
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.3
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.3
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.3
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.3
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.3
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.3
	org.kitesdk	kite-tools	1.0.0-cdh5.4.3
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.3
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.3
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.3
	org.apache.solr	solr-cell	4.10.3-cdh5.4.3
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.3
	org.apache.solr	solr-core	4.10.3-cdh5.4.3

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.3
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.3
	org.apache.solr	solr-langid	4.10.3-cdh5.4.3
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.3
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.3
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.3
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.3
	org.apache.solr	solr-uima	4.10.3-cdh5.4.3
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.3
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.3
	com.cloudera.search	search-mr	1.0.0-cdh5.4.3
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.3
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.3
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.3
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.3
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.3
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.3
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.3
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.3
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.3
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.3
	com.ngdata	hbase-sep-api	1.5-cdh5.4.3
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.3
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.3
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.3
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.3

CDH 5.4.2

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.2
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.4.2
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.4.2
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.4.2
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.2
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.2
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.2
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.2
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.2
	org.apache.hive	hive-ant	1.1.0-cdh5.4.2
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.2
	org.apache.hive	hive-cli	1.1.0-cdh5.4.2
	org.apache.hive	hive-common	1.1.0-cdh5.4.2
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.2
	org.apache.hive	hive-exec	1.1.0-cdh5.4.2
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.2
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.2
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.2
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.2
	org.apache.hive	hive-serde	1.1.0-cdh5.4.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-service	1.1.0-cdh5.4.2
	org.apache.hive	hive-shims	1.1.0-cdh5.4.2
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.2
	org.apache.hive	spark-client	1.1.0-cdh5.4.2
	org.apache.hcatalog	hive-hcatalog-core	1.1.0-cdh5.4.2
	org.apache.hcatalog	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.2
	org.apache.hcatalog	hive-hcatalog-server-extensions	1.1.0-cdh5.4.2
	org.apache.hcatalog	hive-hcatalog-streaming	1.1.0-cdh5.4.2
	org.apache.hcatalog	hive-webhcat	1.1.0-cdh5.4.2
	org.apache.hcatalog	hive-webhcat-java-client	1.1.0-cdh5.4.2
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.4.2
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.4.2
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.4.2
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.2
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.2
	org.apache.pig	piggybank	0.12.0-cdh5.4.2

Project	groupId	artifactId	version
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.2
	org.apache.pig	pigunit	0.12.0-cdh5.4.2
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.2
	org.apache.flume	flume-tools	1.5.0-cdh5.4.2
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.2
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.2
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.2
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.2
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.2
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.2
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.2
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.2
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.2
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.2
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.2
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.2
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.2
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.2
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.2
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.2
	org.apache.oozie	oozie-mini	4.1.0-cdh5.4.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.2
	org.apache.mahout	mahout-core	0.9-cdh5.4.2
	org.apache.mahout	mahout-examples	0.9-cdh5.4.2
	org.apache.mahout	mahout-integration	0.9-cdh5.4.2
	org.apache.mahout	mahout-math	0.9-cdh5.4.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.2
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.4.2
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.2

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.2
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.2
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.2
	com.twitter	parquet-cascading	1.5.0-cdh5.4.2
	com.twitter	parquet-column	1.5.0-cdh5.4.2
	com.twitter	parquet-common	1.5.0-cdh5.4.2
	com.twitter	parquet-encoding	1.5.0-cdh5.4.2
	com.twitter	parquet-format	2.1.0-cdh5.4.2
	com.twitter	parquet-generator	1.5.0-cdh5.4.2
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.2
	com.twitter	parquet-jackson	1.5.0-cdh5.4.2
	com.twitter	parquet-pig	1.5.0-cdh5.4.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.2

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.2
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.2
	com.twitter	parquet-thrift	1.5.0-cdh5.4.2
	com.twitter	parquet-tools	1.5.0-cdh5.4.2
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.2
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.2
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.2

Project	groupId	artifactId	version
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.2
	org.apache.avro	avro-tools	1.7.6-cdh5.4.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.2
	org.apache.avro	trevni-core	1.7.6-cdh5.4.2
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.2
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.2
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.2
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.2
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.2
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.2
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.2
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.2
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.2
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.2
	org.kitesdk	kite-tools	1.0.0-cdh5.4.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.2
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.2
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.2
	org.apache.solr	solr-cell	4.10.3-cdh5.4.2
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.2
	org.apache.solr	solr-core	4.10.3-cdh5.4.2
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.2
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.2
	org.apache.solr	solr-langid	4.10.3-cdh5.4.2
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.2
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.2
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.2
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.2
	org.apache.solr	solr-uima	4.10.3-cdh5.4.2
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.2
	com.cloudera.search	search-mr	1.0.0-cdh5.4.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.2

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.2
	com.ngdata	hbase-sep-api	1.5-cdh5.4.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.2
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.2

CDH 5.4.1

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.1
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.4.1
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.4.1
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.4.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.1
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.1
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.1
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.1
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.1
	org.apache.hive	hive-ant	1.1.0-cdh5.4.1
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.1
	org.apache.hive	hive-cli	1.1.0-cdh5.4.1
	org.apache.hive	hive-common	1.1.0-cdh5.4.1
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.1
	org.apache.hive	hive-exec	1.1.0-cdh5.4.1
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.1
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.1
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.1
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.1
	org.apache.hive	hive-serde	1.1.0-cdh5.4.1
	org.apache.hive	hive-service	1.1.0-cdh5.4.1
	org.apache.hive	hive-shims	1.1.0-cdh5.4.1
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.1
	org.apache.hive	spark-client	1.1.0-cdh5.4.1
	org.apache.hive	hive-hcatalog	1.1.0-cdh5.4.1
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.4.1
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.1
	org.apache.hive	hive-hcatalog-server-extensions	1.1.0-cdh5.4.1
	org.apache.hive	hive-hcatalog-streaming	1.1.0-cdh5.4.1
	org.apache.hive	hive-webhcat	1.1.0-cdh5.4.1
	org.apache.hive	hive-webhcat-java-client	1.1.0-cdh5.4.1
	org.apache.hive	hive-shims-0.23	1.1.0-cdh5.4.1
	org.apache.hive	hive-shims-common	1.1.0-cdh5.4.1
	org.apache.hive	hive-shims-scheduler	1.1.0-cdh5.4.1
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.1

Project	groupId	artifactId	version
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.1
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.1
	org.apache.pig	piggybank	0.12.0-cdh5.4.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.1
	org.apache.pig	pigunit	0.12.0-cdh5.4.1
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.1
	org.apache.flume	flume-tools	1.5.0-cdh5.4.1
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.1
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.1
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.1
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.1
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.1
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.1
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.1
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.1
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.1
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.1
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.1
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.1
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.1
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.1
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.1

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.1
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.1
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.4.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.1
	org.apache.mahout	mahout-core	0.9-cdh5.4.1
	org.apache.mahout	mahout-examples	0.9-cdh5.4.1
	org.apache.mahout	mahout-integration	0.9-cdh5.4.1
	org.apache.mahout	mahout-math	0.9-cdh5.4.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.1
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.4.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.1
	org.apache.sqoop	test	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.4.1
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.4.1
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.1

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.1
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.1
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.1
	com.twitter	parquet-cascading	1.5.0-cdh5.4.1
	com.twitter	parquet-column	1.5.0-cdh5.4.1
	com.twitter	parquet-common	1.5.0-cdh5.4.1
	com.twitter	parquet-encoding	1.5.0-cdh5.4.1
	com.twitter	parquet-format	2.1.0-cdh5.4.1
	com.twitter	parquet-generator	1.5.0-cdh5.4.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.1
	com.twitter	parquet-jackson	1.5.0-cdh5.4.1
	com.twitter	parquet-pig	1.5.0-cdh5.4.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.1
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.1

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.1
	com.twitter	parquet-thrift	1.5.0-cdh5.4.1
	com.twitter	parquet-tools	1.5.0-cdh5.4.1
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.1
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.1
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.1

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.1
	org.apache.avro	avro-tools	1.7.6-cdh5.4.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.1
	org.apache.avro	trevni-core	1.7.6-cdh5.4.1
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.1
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.1
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.1
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.1
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.1
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.1
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.1
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.1
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.1

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.1
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.1
	org.kitesdk	kite-tools	1.0.0-cdh5.4.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.1

Project	groupId	artifactId	version
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.1
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.1
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.1
	org.apache.solr	solr-cell	4.10.3-cdh5.4.1
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.1
	org.apache.solr	solr-core	4.10.3-cdh5.4.1
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.1
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.1
	org.apache.solr	solr-langid	4.10.3-cdh5.4.1
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.1
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.1
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.1
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.1
	org.apache.solr	solr-uima	4.10.3-cdh5.4.1
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.1
	com.cloudera.search	search-mr	1.0.0-cdh5.4.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.1

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.1
	com.ngdata	hbase-sep-api	1.5-cdh5.4.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.1
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.1

CDH 5.4.0

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.4.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-mapreduce-client-nativevtask	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.4.0
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.4.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.4.0
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.4.0
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.4.0
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.4.0
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.4.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.4.0
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.4.0
Apache Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.4.0
	org.apache.hive	hive-ant	1.1.0-cdh5.4.0
	org.apache.hive	hive-beeline	1.1.0-cdh5.4.0
	org.apache.hive	hive-cli	1.1.0-cdh5.4.0
	org.apache.hive	hive-common	1.1.0-cdh5.4.0
	org.apache.hive	hive-contrib	1.1.0-cdh5.4.0
	org.apache.hive	hive-exec	1.1.0-cdh5.4.0
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.4.0
	org.apache.hive	hive-hwi	1.1.0-cdh5.4.0
	org.apache.hive	hive-jdbc	1.1.0-cdh5.4.0
	org.apache.hive	hive-metastore	1.1.0-cdh5.4.0
	org.apache.hive	hive-serde	1.1.0-cdh5.4.0
	org.apache.hive	hive-service	1.1.0-cdh5.4.0
	org.apache.hive	hive-shims	1.1.0-cdh5.4.0
	org.apache.hive	hive-testutils	1.1.0-cdh5.4.0
	org.apache.hive	spark-client	1.1.0-cdh5.4.0
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.4.0
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.4.0
	org.apache.hive	hive-hcatalog-server-extensions	1.1.0-cdh5.4.0
	org.apache.hive	hive-hcatalog-streaming	1.1.0-cdh5.4.0
	org.apache.hive	hive-webhcat	1.1.0-cdh5.4.0
	org.apache.hive	hive-webhcat-java-client	1.1.0-cdh5.4.0
	org.apache.hiveshims	hive-shims-0.23	1.1.0-cdh5.4.0
	org.apache.hiveshims	hive-shims-common	1.1.0-cdh5.4.0
	org.apache.hiveshims	hive-shims-scheduler	1.1.0-cdh5.4.0
Apache HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-client	1.0.0-cdh5.4.0

Project	groupId	artifactId	version
	org.apache.hbase	hbase-common	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-examples	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-it	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-rest	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-server	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-shell	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.4.0
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.4.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.4.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.4.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.4.0
	org.apache.pig	piggybank	0.12.0-cdh5.4.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.4.0
	org.apache.pig	pigunit	0.12.0-cdh5.4.0
Apache Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-core	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-node	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.4.0
	org.apache.flume	flume-tools	1.5.0-cdh5.4.0
	org.apache.flume	flume-file-channel	1.5.0-cdh5.4.0
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.4.0
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.4.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.4.0
	org.apache.flume	flume-avro-source	1.5.0-cdh5.4.0
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.4.0
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.4.0
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.4.0
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.4.0
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.4.0
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.4.0
	org.apache.flume	flume-jms-source	1.5.0-cdh5.4.0
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.4.0
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.4.0
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.4.0
Apache Oozie	org.apache.oozie	oozie-client	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-core	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-examples	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-hadoop	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-hadoop-distcp	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-hadoop-test	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-hadoop-utils	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-hbase	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-hcatalog	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.4.0

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.4.0
	org.apache.oozie	oozie-tools	4.1.0-cdh5.4.0
	org.apache.oozie.test	oozie-mini	4.1.0-cdh5.4.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.4.0
	org.apache.mahout	mahout-core	0.9-cdh5.4.0
	org.apache.mahout	mahout-examples	0.9-cdh5.4.0
	org.apache.mahout	mahout-integration	0.9-cdh5.4.0
	org.apache.mahout	mahout-math	0.9-cdh5.4.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.4.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.4.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.4.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.4.0
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.4.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-client	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-common	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-common-test	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-core	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-docs	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-security	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-shell	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-tomcat	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-tools	1.99.5-cdh5.4.0
	org.apache.sqoop	test	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-connector-hdfs	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-connector-kafka	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-connector-kite	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-execution-mapreduce	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-repository-common	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-repository-derby	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-repository-postgresql	1.99.5-cdh5.4.0
	org.apache.sqoop	sqoop-submission-mapreduce	1.99.5-cdh5.4.0
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hbase-indexer	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-core-model-indexer	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.4.0

Project	groupId	artifactId	version
	org.apache.sentry	sentry-dist	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-policy-indexer	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.4.0
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.4.0
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.4.0
	com.twitter	parquet-cascading	1.5.0-cdh5.4.0
	com.twitter	parquet-column	1.5.0-cdh5.4.0
	com.twitter	parquet-common	1.5.0-cdh5.4.0
	com.twitter	parquet-encoding	1.5.0-cdh5.4.0
	com.twitter	parquet-format	2.1.0-cdh5.4.0
	com.twitter	parquet-generator	1.5.0-cdh5.4.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.4.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.4.0
	com.twitter	parquet-jackson	1.5.0-cdh5.4.0
	com.twitter	parquet-pig	1.5.0-cdh5.4.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.4.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.4.0
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.4.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.4.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.4.0

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-thrift	1.5.0-cdh5.4.0
	com.twitter	parquet-tools	1.5.0-cdh5.4.0
Llama	com.doudera.llama	llama	1.0.0-cdh5.4.0
Apache Spark	org.apache.spark	spark-assembly_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-bagel_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-catalyst_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-core_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-graphx_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-hive_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-mllib_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-network-common_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-network-shuffle_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-network-yarn_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-repl_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-sql_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-streaming-flume_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-streaming-kafka-assembly_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-streaming-kafka_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-streaming-twitter_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-streaming_2.10	1.3.0-cdh5.4.0
	org.apache.spark	spark-yarn_2.10	1.3.0-cdh5.4.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.4.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.4.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.4.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.4.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.4.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.4.0

Project	groupId	artifactId	version
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.4.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.4.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.4.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.4.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.4.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.4.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.4.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.4.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.4.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.4.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.4.0
	org.apache.avro	avro-tools	1.7.6-cdh5.4.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.4.0
	org.apache.avro	trevni-core	1.7.6-cdh5.4.0
Kite SDK	org.kitesdk	kite-data-core	1.0.0-cdh5.4.0
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.4.0
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.4.0
	org.kitesdk	kite-data-hive	1.0.0-cdh5.4.0
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.4.0
	org.kitesdk	kite-data-spark	1.0.0-cdh5.4.0
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.4.0
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.4.0
	org.kitesdk	kite-minicluster	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.4.0

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.4.0
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.4.0
	org.kitesdk	kite-tools	1.0.0-cdh5.4.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-analyzers-icu	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-analyzers-morfologik	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-analyzers-phonetic	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-analyzers-smartcn	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-analyzers-stempel	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-analyzers-uima	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-benchmark	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-classification	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-codecs	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-core	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-demo	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-expressions	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-facet	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-grouping	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-highlighter	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-join	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-memory	4.10.3-cdh5.4.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-misc	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-queries	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-queryparser	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-replicator	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-sandbox	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-spatial	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-suggest	4.10.3-cdh5.4.0
	org.apache.lucene	lucene-test-framework	4.10.3-cdh5.4.0
	org.apache.solr	solr-analysis-extras	4.10.3-cdh5.4.0
	org.apache.solr	solr-cell	4.10.3-cdh5.4.0
	org.apache.solr	solr-clustering	4.10.3-cdh5.4.0
	org.apache.solr	solr-core	4.10.3-cdh5.4.0
	org.apache.solr	solr-dataimporthandler	4.10.3-cdh5.4.0
	org.apache.solr	solr-dataimporthandler-extras	4.10.3-cdh5.4.0
	org.apache.solr	solr-langid	4.10.3-cdh5.4.0
	org.apache.solr	solr-security-util	4.10.3-cdh5.4.0
	org.apache.solr	solr-sentry-handlers	4.10.3-cdh5.4.0
	org.apache.solr	solr-solrj	4.10.3-cdh5.4.0
	org.apache.solr	solr-test-framework	4.10.3-cdh5.4.0
	org.apache.solr	solr-uima	4.10.3-cdh5.4.0
	org.apache.solr	solr-velocity	4.10.3-cdh5.4.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.4.0
	com.cloudera.search	search-mr	1.0.0-cdh5.4.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.4.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.4.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.4.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.4.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.4.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.4.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.4.0

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.4.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.4.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.4.0
	com.ngdata	hbase-sep-api	1.5-cdh5.4.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.4.0
	com.ngdata	hbase-sep-impl	1.5-cdh5.4.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.4.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.4.0

Maven Artifacts for CDH 5.3.x Releases

CDH 5.3.10

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-aws	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-azure	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.3.10

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.3.10
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.10
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	250mr1-cdh5.3.10
	org.apache.hadoop	hadoop-core	250mr1-cdh5.3.10
	org.apache.hadoop	hadoop-examples	250mr1-cdh5.3.10
	org.apache.hadoop	hadoop-minicluster	250mr1-cdh5.3.10
	org.apache.hadoop	hadoop-streaming	250mr1-cdh5.3.10
	org.apache.hadoop	hadoop-test	250mr1-cdh5.3.10
	org.apache.hadoop	hadoop-tools	250mr1-cdh5.3.10
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.3.10
	org.apache.hive	hive-beeline	0.13.1-cdh5.3.10

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-cli	0.13.1-cdh5.3.10
	org.apache.hive	hive-common	0.13.1-cdh5.3.10
	org.apache.hive	hive-contrib	0.13.1-cdh5.3.10
	org.apache.hive	hive-exec	0.13.1-cdh5.3.10
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.3.10
	org.apache.hive	hive-hwi	0.13.1-cdh5.3.10
	org.apache.hive	hive-jdbc	0.13.1-cdh5.3.10
	org.apache.hive	hive-metastore	0.13.1-cdh5.3.10
	org.apache.hive	hive-serde	0.13.1-cdh5.3.10
	org.apache.hive	hive-service	0.13.1-cdh5.3.10
	org.apache.hive	hive-shims	0.13.1-cdh5.3.10
	org.apache.hive	hive-testutils	0.13.1-cdh5.3.10
	org.apache.hcatalog	hive-hcatalog-core	0.13.1-cdh5.3.10
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.10
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.13.1-cdh5.3.10
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.3.10
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.3.10
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.3.10
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.3.10
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.3.10
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.3.10
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.3.10
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.3.10
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-common	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-examples	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-it	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.3.10

Project	groupId	artifactId	version
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-server	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-shell	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.3.10
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.3.10
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.3.10
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.10
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.3.10
	org.apache.pig	piggybank	0.12.0-cdh5.3.10
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.10
	org.apache.pig	pigunit	0.12.0-cdh5.3.10
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.10
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.10
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.10
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.10
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.10
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.10
	org.apache.flume	flume-tools	1.5.0-cdh5.3.10
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.10
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.10
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.10
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.10
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.10
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.10
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.10
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.10
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.10
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.10
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.10

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.10
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.10
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.10
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.10
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.10
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.10
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.10
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-core	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-examples	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-tools	4.0.0-cdh5.3.10
	org.apache.oozie	oozie-mini	4.0.0-cdh5.3.10
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.10
	org.apache.mahout	mahout-core	0.9-cdh5.3.10
	org.apache.mahout	mahout-examples	0.9-cdh5.3.10
	org.apache.mahout	mahout-integration	0.9-cdh5.3.10

Project	groupId	artifactId	version
	org.apache.mahout	mahout-math	0.9-cdh5.3.10
	org.apache.mahout	mahout-math-scala	0.9-cdh5.3.10
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.10
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.10
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.10
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.10

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.10
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.10
Apache Sentry	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.10
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.10
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.10
	com.twitter	parquet-cascading	1.5.0-cdh5.3.10
	com.twitter	parquet-column	1.5.0-cdh5.3.10
	com.twitter	parquet-common	1.5.0-cdh5.3.10
	com.twitter	parquet-encoding	1.5.0-cdh5.3.10
	com.twitter	parquet-format	2.1.0-cdh5.3.10

Project	groupId	artifactId	version
	com.twitter	parquet-generator	1.5.0-cdh5.3.10
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.10
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.10
	com.twitter	parquet-jackson	1.5.0-cdh5.3.10
	com.twitter	parquet-pig	1.5.0-cdh5.3.10
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.10
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.10
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.3.10
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.3.10
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.10
	com.twitter	parquet-thrift	1.5.0-cdh5.3.10
	com.twitter	parquet-tools	1.5.0-cdh5.3.10
Llama	com.doudera.llama	llama	1.0.0-cdh5.3.10
Apache Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.10
	org.apache.spark	spark-streaming_2.10	1.2.0-cdh5.3.10

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-yarn_2.10	1.2.0-cdh5.3.10
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.3.10
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.3.10
	org.apache.crunch	crunch-core	0.11.0-cdh5.3.10
	org.apache.crunch	crunch-examples	0.11.0-cdh5.3.10
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.3.10
	org.apache.crunch	crunch-hive	0.11.0-cdh5.3.10
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.3.10
	org.apache.crunch	crunch-spark	0.11.0-cdh5.3.10
	org.apache.crunch	crunch-test	0.11.0-cdh5.3.10
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.3.10
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.10
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.10
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.10
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.10
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.10
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.10
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.10
	org.apache.avro	avro-tools	1.7.6-cdh5.3.10
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.10
	org.apache.avro	trevni-core	1.7.6-cdh5.3.10
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.3.10
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.3.10
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.3.10
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.10
	org.kitesdk	kite-data-hive	0.15.0-cdh5.3.10
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.3.10
	org.kitesdk	kite-data-spark	0.15.0-cdh5.3.10
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.3.10
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.3.10

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.3.10
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.3.10
	org.kitesdk	kite-tools	0.15.0-cdh5.3.10
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.10

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.10
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.10
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.10
	org.apache.solr	solr-cell	4.4.0-cdh5.3.10
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.10
	org.apache.solr	solr-core	4.4.0-cdh5.3.10
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.10
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.10
	org.apache.solr	solr-langid	4.4.0-cdh5.3.10
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.10
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.10
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.10
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.10
	org.apache.solr	solr-uima	4.4.0-cdh5.3.10
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.10
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.3.10
	com.cloudera.search	search-mr	1.0.0-cdh5.3.10
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.10

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.10
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.10
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.10
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.10
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.10
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.10
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.10
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.10
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.10
	com.ngdata	hbase-sep-api	1.5-cdh5.3.10
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.10
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.10
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.10
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.10

CDH 5.3.9

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-aws	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-azure	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.3.9

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.9
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-core	2.5.0-cdh5.3.9
	org.apache.hadoop	hadoop-examples	2.5.0-cdh5.3.9

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-minicluster	2.5.0-mr1-cdh5.3.9
	org.apache.hadoop	hadoop-streaming	2.5.0-mr1-cdh5.3.9
	org.apache.hadoop	hadoop-test	2.5.0-mr1-cdh5.3.9
	org.apache.hadoop	hadoop-tools	2.5.0-mr1-cdh5.3.9
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.3.9
	org.apache.hive	hive-beeline	0.13.1-cdh5.3.9
	org.apache.hive	hive-cli	0.13.1-cdh5.3.9
	org.apache.hive	hive-common	0.13.1-cdh5.3.9
	org.apache.hive	hive-contrib	0.13.1-cdh5.3.9
	org.apache.hive	hive-exec	0.13.1-cdh5.3.9
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.3.9
	org.apache.hive	hive-hwi	0.13.1-cdh5.3.9
	org.apache.hive	hive-jdbc	0.13.1-cdh5.3.9
	org.apache.hive	hive-metastore	0.13.1-cdh5.3.9
	org.apache.hive	hive-serde	0.13.1-cdh5.3.9
	org.apache.hive	hive-service	0.13.1-cdh5.3.9
	org.apache.hive	hive-shims	0.13.1-cdh5.3.9
	org.apache.hive	hive-testutils	0.13.1-cdh5.3.9
	org.apache.hive	hive-hcatalog-core	0.13.1-cdh5.3.9
	org.apache.hive	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.9
	org.apache.hive	hive-hcatalog-pig-adapter	0.13.1-cdh5.3.9
	org.apache.hive	hive-hcatalog-server-extensions	0.13.1-cdh5.3.9
	org.apache.hive	hive-hcatalog-streaming	0.13.1-cdh5.3.9
	org.apache.hive	hive-webhcat	0.13.1-cdh5.3.9
	org.apache.hive	hive-webhcat-java-client	0.13.1-cdh5.3.9
	org.apache.hive	hive-shims-0.23	0.13.1-cdh5.3.9
	org.apache.hive	hive-shims-common	0.13.1-cdh5.3.9
	org.apache.hive	hive-shims-common-secure	0.13.1-cdh5.3.9
	org.apache.hive	hive-shims-scheduler	0.13.1-cdh5.3.9
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.3.9

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-common	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-examples	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-it	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-server	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-shell	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.3.9
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.3.9
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.3.9
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.9
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.3.9
	org.apache.pig	piggybank	0.12.0-cdh5.3.9
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.9
	org.apache.pig	pigunit	0.12.0-cdh5.3.9
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.9
	org.apache.flume	flume-tools	1.5.0-cdh5.3.9
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.9
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.9
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.9
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.9

Project	groupId	artifactId	version
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.9
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.9
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.9
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.9
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.9
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.9
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.9
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.9
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.9
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.9
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-core	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-examples	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.3.9
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.3.9

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-tools	4.0.0-cdh5.3.9
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.3.9
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.9
	org.apache.mahout	mahout-core	0.9-cdh5.3.9
	org.apache.mahout	mahout-examples	0.9-cdh5.3.9
	org.apache.mahout	mahout-integration	0.9-cdh5.3.9
	org.apache.mahout	mahout-math	0.9-cdh5.3.9
	org.apache.mahout	mahout-math-scala	0.9-cdh5.3.9
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.9
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.9
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.9
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.9

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.9
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.9
Apache Sentry	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.9
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.9

Version and Download Information

Project	groupId	artifactId	version
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.9
	com.twitter	parquet-cascading	1.5.0-cdh5.3.9
	com.twitter	parquet-column	1.5.0-cdh5.3.9
	com.twitter	parquet-common	1.5.0-cdh5.3.9
	com.twitter	parquet-encoding	1.5.0-cdh5.3.9
	com.twitter	parquet-format	2.1.0-cdh5.3.9
	com.twitter	parquet-generator	1.5.0-cdh5.3.9
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.9
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.9
	com.twitter	parquet-jackson	1.5.0-cdh5.3.9
	com.twitter	parquet-pig	1.5.0-cdh5.3.9
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.9
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.9
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.3.9
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.3.9
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.9
	com.twitter	parquet-thrift	1.5.0-cdh5.3.9
	com.twitter	parquet-tools	1.5.0-cdh5.3.9
Llama	com.doudera.llama	llama	1.0.0-cdh5.3.9
Apache Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.9

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-streaming_2.10	1.2.0-cdh5.3.9
	org.apache.spark	spark-yarn_2.10	1.2.0-cdh5.3.9
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.3.9
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.3.9
	org.apache.crunch	crunch-core	0.11.0-cdh5.3.9
	org.apache.crunch	crunch-examples	0.11.0-cdh5.3.9
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.3.9
	org.apache.crunch	crunch-hive	0.11.0-cdh5.3.9
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.3.9
	org.apache.crunch	crunch-spark	0.11.0-cdh5.3.9
	org.apache.crunch	crunch-test	0.11.0-cdh5.3.9
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.3.9
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.9
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.9
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.9
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.9
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.9
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.9
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.9
	org.apache.avro	avro-tools	1.7.6-cdh5.3.9
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.9
	org.apache.avro	trevni-core	1.7.6-cdh5.3.9
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.3.9
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.3.9
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.3.9

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.9
	org.kitesdk	kite-data-hive	0.15.0-cdh5.3.9
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.3.9
	org.kitesdk	kite-data-spark	0.15.0-cdh5.3.9
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.3.9
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.3.9
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.3.9
	org.kitesdk	kite-tools	0.15.0-cdh5.3.9
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.9

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.9
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.9
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.9
	org.apache.solr	solr-cell	4.4.0-cdh5.3.9
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.9
	org.apache.solr	solr-core	4.4.0-cdh5.3.9
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.9
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.9
	org.apache.solr	solr-langid	4.4.0-cdh5.3.9
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.9
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.9
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.9

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.9
	org.apache.solr	solr-uima	4.4.0-cdh5.3.9
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.9
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.3.9
	com.cloudera.search	search-mr	1.0.0-cdh5.3.9
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.9
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.9
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.9
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.9
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.9
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.9
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.9
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.9
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.9
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.9
	com.ngdata	hbase-sep-api	1.5-cdh5.3.9
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.9
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.9
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.9
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.9

CDH 5.3.8

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-aws	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-azure	2.5.0-cdh5.3.8

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.3.8

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.3.8
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.8
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	250-mr1-cdh5.3.8
	org.apache.hadoop	hadoop-core	250-mr1-cdh5.3.8
	org.apache.hadoop	hadoop-examples	250-mr1-cdh5.3.8
	org.apache.hadoop	hadoop-minicluster	250-mr1-cdh5.3.8
	org.apache.hadoop	hadoop-streaming	250-mr1-cdh5.3.8
	org.apache.hadoop	hadoop-test	250-mr1-cdh5.3.8
	org.apache.hadoop	hadoop-tools	250-mr1-cdh5.3.8
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.3.8
	org.apache.hive	hive-beeline	0.13.1-cdh5.3.8
	org.apache.hive	hive-cli	0.13.1-cdh5.3.8
	org.apache.hive	hive-common	0.13.1-cdh5.3.8
	org.apache.hive	hive-contrib	0.13.1-cdh5.3.8
	org.apache.hive	hive-exec	0.13.1-cdh5.3.8
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.3.8
	org.apache.hive	hive-hwi	0.13.1-cdh5.3.8
	org.apache.hive	hive-jdbc	0.13.1-cdh5.3.8
	org.apache.hive	hive-metastore	0.13.1-cdh5.3.8
	org.apache.hive	hive-serde	0.13.1-cdh5.3.8
	org.apache.hive	hive-service	0.13.1-cdh5.3.8
	org.apache.hive	hive-shims	0.13.1-cdh5.3.8
	org.apache.hive	hive-testutils	0.13.1-cdh5.3.8
	org.apache.hive	hive-hcatalog-core	0.13.1-cdh5.3.8
	org.apache.hive	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.8
	org.apache.hive	hive-hcatalog-pig-adapter	0.13.1-cdh5.3.8
	org.apache.hive	hive-hcatalog-server-extensions	0.13.1-cdh5.3.8
	org.apache.hive	hive-hcatalog-streaming	0.13.1-cdh5.3.8

Project	groupId	artifactId	version
	org.apache.hadoop	hive-webhcat	0.13.1-cdh5.3.8
	org.apache.hadoop	hive-webhcat-java-client	0.13.1-cdh5.3.8
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.3.8
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.3.8
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.3.8
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.3.8
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-common	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-examples	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-it	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-server	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-shell	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.3.8
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.3.8
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.3.8
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.8
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.3.8
	org.apache.pig	piggybank	0.12.0-cdh5.3.8
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.8
	org.apache.pig	pigunit	0.12.0-cdh5.3.8
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.8
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.8
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.8
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.8
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.8

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.8
	org.apache.flume	flume-tools	1.5.0-cdh5.3.8
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.8
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.8
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.8
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.8
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.8
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.8
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.8
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.8
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.8
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.8
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.8
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.8
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.8
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.8
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.8
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.8
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.8
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.8
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-core	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-examples	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.3.8

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.3.8
	org.apache.oozie	oozie-tools	4.0.0-cdh5.3.8
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.3.8
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.8
	org.apache.mahout	mahout-core	0.9-cdh5.3.8
	org.apache.mahout	mahout-examples	0.9-cdh5.3.8
	org.apache.mahout	mahout-integration	0.9-cdh5.3.8
	org.apache.mahout	mahout-math	0.9-cdh5.3.8
	org.apache.mahout	mahout-math-scala	0.9-cdh5.3.8
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.8

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.8
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.8
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.8
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.8
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.8
Apache Sentry	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.8

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.8
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.8
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.8
	com.twitter	parquet-cascading	1.5.0-cdh5.3.8
	com.twitter	parquet-column	1.5.0-cdh5.3.8
	com.twitter	parquet-common	1.5.0-cdh5.3.8
	com.twitter	parquet-encoding	1.5.0-cdh5.3.8
	com.twitter	parquet-format	2.1.0-cdh5.3.8
	com.twitter	parquet-generator	1.5.0-cdh5.3.8
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.8
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.8
	com.twitter	parquet-jackson	1.5.0-cdh5.3.8
	com.twitter	parquet-pig	1.5.0-cdh5.3.8
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.8
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.8
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.3.8
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.3.8
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.8
	com.twitter	parquet-thrift	1.5.0-cdh5.3.8
	com.twitter	parquet-tools	1.5.0-cdh5.3.8
Llama	com.doudera.llama	llama	1.0.0-cdh5.3.8
Apache Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.8

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-streaming_2.10	1.2.0-cdh5.3.8
	org.apache.spark	spark-yarn_2.10	1.2.0-cdh5.3.8
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.3.8
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.3.8
	org.apache.crunch	crunch-core	0.11.0-cdh5.3.8
	org.apache.crunch	crunch-examples	0.11.0-cdh5.3.8
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.3.8
	org.apache.crunch	crunch-hive	0.11.0-cdh5.3.8
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.3.8
	org.apache.crunch	crunch-spark	0.11.0-cdh5.3.8
	org.apache.crunch	crunch-test	0.11.0-cdh5.3.8
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.3.8
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.8
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.8
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.8
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.8
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.8
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.8

Project	groupId	artifactId	version
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.8
	org.apache.avro	avro-tools	1.7.6-cdh5.3.8
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.8
	org.apache.avro	trevni-core	1.7.6-cdh5.3.8
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.3.8
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.3.8
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.3.8
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.8
	org.kitesdk	kite-data-hive	0.15.0-cdh5.3.8
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.3.8
	org.kitesdk	kite-data-spark	0.15.0-cdh5.3.8
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.3.8
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.3.8
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.3.8

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-tools	0.15.0-cdh5.3.8
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.8
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.8
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.8
	org.apache.solr	solr-cell	4.4.0-cdh5.3.8
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.8

Project	groupId	artifactId	version
	org.apache.solr	solr-core	4.4.0-cdh5.3.8
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.8
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.8
	org.apache.solr	solr-langid	4.4.0-cdh5.3.8
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.8
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.8
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.8
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.8
	org.apache.solr	solr-uima	4.4.0-cdh5.3.8
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.8
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.3.8
	com.cloudera.search	search-mr	1.0.0-cdh5.3.8
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.8
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.8
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.8
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.8
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.8
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.8
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.8
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.8
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.8
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.8
	com.ngdata	hbase-sep-api	1.5-cdh5.3.8
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.8
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.8
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.8
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.8

CDH 5.3.6

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Version and Download Information

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.6
	org.apache.hadoop	hadoop-annotations	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-ant	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-archives	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-assemblies	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-auth	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-aws	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-azure	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-common	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-datajoin	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-distcp	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-extras	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-gridmix	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-hdfs	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-hdfs-nfs	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-mapreduce-client-app	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-mapreduce-client-common	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-mapreduce-client-core	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-mapreduce-examples	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-maven-plugins	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-minikdc	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-nfs	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-openstack	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-rumen	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-sls	2.6.0-cdh5.3.6

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-api	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-client	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-common	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-registry	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-server-common	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-server-tests	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.6.0-cdh5.3.6
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.6
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.6.0-cdh5.3.6
Hadoop MRv1	org.apache.hadoop	hadoop-client	2.6.0-mr1-cdh5.3.6
	org.apache.hadoop	hadoop-core	2.6.0-mr1-cdh5.3.6
	org.apache.hadoop	hadoop-examples	2.6.0-mr1-cdh5.3.6
	org.apache.hadoop	hadoop-minicluster	2.6.0-mr1-cdh5.3.6
	org.apache.hadoop	hadoop-streaming	2.6.0-mr1-cdh5.3.6
	org.apache.hadoop	hadoop-test	2.6.0-mr1-cdh5.3.6
	org.apache.hadoop	hadoop-tools	2.6.0-mr1-cdh5.3.6
Hive	org.apache.hive	hive-accumulo-handler	1.1.0-cdh5.3.6
	org.apache.hive	hive-ant	1.1.0-cdh5.3.6
	org.apache.hive	hive-beeline	1.1.0-cdh5.3.6
	org.apache.hive	hive-cli	1.1.0-cdh5.3.6
	org.apache.hive	hive-common	1.1.0-cdh5.3.6
	org.apache.hive	hive-contrib	1.1.0-cdh5.3.6
	org.apache.hive	hive-exec	1.1.0-cdh5.3.6
	org.apache.hive	hive-hbase-handler	1.1.0-cdh5.3.6
	org.apache.hive	hive-hwi	1.1.0-cdh5.3.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-jdbc	1.1.0-cdh5.3.6
	org.apache.hive	hive-metastore	1.1.0-cdh5.3.6
	org.apache.hive	hive-serde	1.1.0-cdh5.3.6
	org.apache.hive	hive-service	1.1.0-cdh5.3.6
	org.apache.hive	hive-shims	1.1.0-cdh5.3.6
	org.apache.hive	hive-testutils	1.1.0-cdh5.3.6
	org.apache.hive	spark-client	1.1.0-cdh5.3.6
	org.apache.hive	hive-hcatalog	0.13.1-cdh5.3.6
	org.apache.hive	hive-hcatalog-core	1.1.0-cdh5.3.6
	org.apache.hive	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.6
	org.apache.hive	hive-hcatalog-pig-adapter	1.1.0-cdh5.3.6
	org.apache.hive	hive-hcatalog-server-extensions	1.1.0-cdh5.3.6
	org.apache.hive	hive-hcatalog-streaming	1.1.0-cdh5.3.6
	org.apache.hive	hive-webhcat	1.1.0-cdh5.3.6
	org.apache.hive	hive-webhcat-java-client	1.1.0-cdh5.3.6
	org.apache.hive	hive-shims-0.23	0.13.1-cdh5.3.6
	org.apache.hive	hive-shims-0.23	1.1.0-cdh5.3.6
	org.apache.hive	hive-shims-common	1.1.0-cdh5.3.6
	org.apache.hive	hive-shims-common-secure	0.13.1-cdh5.3.6
	org.apache.hive	hive-shims-scheduler	1.1.0-cdh5.3.6
HBase	org.apache.hbase	hbase-annotations	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-checkstyle	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-client	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-common	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-examples	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-hadoop-compat	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-hadoop2-compat	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-it	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-prefix-tree	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-protocol	1.0.0-cdh5.3.6

Project	groupId	artifactId	version
	org.apache.hbase	hbase-rest	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-server	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-shell	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-testing-util	1.0.0-cdh5.3.6
	org.apache.hbase	hbase-thrift	1.0.0-cdh5.3.6
ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.3.6
Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.6
Pig	org.apache.pig	pig	0.12.0-cdh5.3.6
	org.apache.pig	piggybank	0.12.0-cdh5.3.6
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.6
	org.apache.pig	pigunit	0.12.0-cdh5.3.6
Flume 1.x	org.apache.flume	flume-ng-auth	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.6
	org.apache.flume	flume-tools	1.5.0-cdh5.3.6
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.6
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.6
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.6
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.6
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.6
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.6
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.6
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.6
	org.apache.flume	flume-hive-sink	1.5.0-cdh5.3.6
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.6
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.6
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.6
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.6
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.6
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.6
Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.6
	org.apache.oozie	oozie-client	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-core	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-examples	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-hadoop	2014156-cdh5.3.6
	org.apache.oozie	oozie-hadoop-distcp	2014156-cdh5.3.6
	org.apache.oozie	oozie-hadoop-test	2014156-cdh5.3.6
	org.apache.oozie	oozie-hadoop-utils	2014156-cdh5.3.6
	org.apache.oozie	oozie-hbase	10-cdh5.3.6-cdh5.3.6
	org.apache.oozie	oozie-hcatalog	10-cdh5.3.6-cdh5.3.6
	org.apache.oozie	oozie-sharelib-distcp	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-sharelib-hcatalog	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-sharelib-hive	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-sharelib-hive2	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-sharelib-oozie	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-sharelib-pig	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-sharelib-spark	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-sharelib-sqoop	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-sharelib-streaming	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-tools	4.1.0-cdh5.3.6
	org.apache.oozie	oozie-mini	4.0.0-cdh5.3.6
	org.apache.oozie	oozie-mini	4.1.0-cdh5.3.6

Project	groupId	artifactId	version
Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.6
	org.apache.mahout	mahout-core	0.9-cdh5.3.6
	org.apache.mahout	mahout-examples	0.9-cdh5.3.6
	org.apache.mahout	mahout-integration	0.9-cdh5.3.6
	org.apache.mahout	mahout-math	0.9-cdh5.3.6
	org.apache.mahout	mahout-math-scala	0.9-cdh5.3.6
Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.6
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.6
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.6
Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.6
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.6
Sentry	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.6
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.6
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.6
	com.twitter	parquet-cascading	1.5.0-cdh5.3.6

Project	groupId	artifactId	version
	com.twitter	parquet-column	1.5.0-cdh5.3.6
	com.twitter	parquet-common	1.5.0-cdh5.3.6
	com.twitter	parquet-encoding	1.5.0-cdh5.3.6
	com.twitter	parquet-format	2.1.0-cdh5.3.6
	com.twitter	parquet-generator	1.5.0-cdh5.3.6
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.6
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.6
	com.twitter	parquet-jackson	1.5.0-cdh5.3.6
	com.twitter	parquet-pig	1.5.0-cdh5.3.6
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.6
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.6
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.3.6
	com.twitter	parquet-scooqe_2.10	1.5.0-cdh5.3.6
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.6
	com.twitter	parquet-thrift	1.5.0-cdh5.3.6
	com.twitter	parquet-tools	1.5.0-cdh5.3.6
Llama	com.doudera.llama	llama	1.0.0-cdh5.3.6
Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-streaming_2.10	1.2.0-cdh5.3.6
	org.apache.spark	spark-yarn_2.10	1.2.0-cdh5.3.6
Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.3.6
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.3.6
	org.apache.crunch	crunch-core	0.11.0-cdh5.3.6
	org.apache.crunch	crunch-examples	0.11.0-cdh5.3.6
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.3.6
	org.apache.crunch	crunch-hive	0.11.0-cdh5.3.6
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.3.6
	org.apache.crunch	crunch-spark	0.11.0-cdh5.3.6
	org.apache.crunch	crunch-test	0.11.0-cdh5.3.6
Avro	org.apache.avro	avro	1.7.6-cdh5.3.6
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.6
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.6
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.6
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.6
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.6
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.6
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.6
	org.apache.avro	avro-tools	1.7.6-cdh5.3.6
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.6
	org.apache.avro	trevni-core	1.7.6-cdh5.3.6
Kite	org.kitesdk	kite-data-core	0.15.0-cdh5.3.6
	org.kitesdk	kite-data-core	1.0.0-cdh5.3.6
	org.kitesdk	kite-data-crunch	1.0.0-cdh5.3.6
	org.kitesdk	kite-data-hbase	1.0.0-cdh5.3.6
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.6

Project	groupId	artifactId	version
	org.kitesdk	kite-data-hive	1.0.0-cdh5.3.6
	org.kitesdk	kite-data-mapreduce	1.0.0-cdh5.3.6
	org.kitesdk	kite-data-spark	1.0.0-cdh5.3.6
	org.kitesdk	kite-hadoop-compatibility	1.0.0-cdh5.3.6
	org.kitesdk	kite-maven-plugin	1.0.0-cdh5.3.6
	org.kitesdk	kite-minicluster	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-avro	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-core	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-hadoop-core	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-hadoop-rcfile	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-hadoop-sequencefile	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-json	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-maxmind	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-metrics-servlets	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-protobuf	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-saxon	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-solr-cell	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-solr-core	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-tika-core	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-tika-decompress	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-twitter	1.0.0-cdh5.3.6
	org.kitesdk	kite-morphlines-useragent	1.0.0-cdh5.3.6
	org.kitesdk	kite-tools	1.0.0-cdh5.3.6
Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.6
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.6
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.6
	org.apache.solr	solr-cell	4.4.0-cdh5.3.6
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.6
	org.apache.solr	solr-core	4.4.0-cdh5.3.6
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.6
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.6
	org.apache.solr	solr-langid	4.4.0-cdh5.3.6
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.6
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.6
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.6

Project	groupId	artifactId	version
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.6
	org.apache.solr	solr-uima	4.4.0-cdh5.3.6
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.6
Search	com.doudera.search	search-crunch	1.0.0-cdh5.3.6
	com.doudera.search	search-mr	1.0.0-cdh5.3.6
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.6
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.6
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.6
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.6
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.6
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.6
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.6
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.6
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.6
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.6
	com.ngdata	hbase-sep-api	1.5-cdh5.3.6
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.6
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.6
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.6
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.6

CDH 5.3.5

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-aws	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-azure	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.3.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.3.5

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.3.5
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.5
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	250-mr1-cdh5.3.5
	org.apache.hadoop	hadoop-core	250-mr1-cdh5.3.5
	org.apache.hadoop	hadoop-examples	250-mr1-cdh5.3.5
	org.apache.hadoop	hadoop-minicluster	250-mr1-cdh5.3.5
	org.apache.hadoop	hadoop-streaming	250-mr1-cdh5.3.5
	org.apache.hadoop	hadoop-test	250-mr1-cdh5.3.5
	org.apache.hadoop	hadoop-tools	250-mr1-cdh5.3.5
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.3.5
	org.apache.hive	hive-beeline	0.13.1-cdh5.3.5
	org.apache.hive	hive-cli	0.13.1-cdh5.3.5
	org.apache.hive	hive-common	0.13.1-cdh5.3.5
	org.apache.hive	hive-contrib	0.13.1-cdh5.3.5
	org.apache.hive	hive-exec	0.13.1-cdh5.3.5
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.3.5
	org.apache.hive	hive-hwi	0.13.1-cdh5.3.5
	org.apache.hive	hive-jdbc	0.13.1-cdh5.3.5
	org.apache.hive	hive-metastore	0.13.1-cdh5.3.5
	org.apache.hive	hive-serde	0.13.1-cdh5.3.5
	org.apache.hive	hive-service	0.13.1-cdh5.3.5
	org.apache.hive	hive-shims	0.13.1-cdh5.3.5
	org.apache.hive	hive-testutils	0.13.1-cdh5.3.5
	org.apache.hive	hive-hcatalog-core	0.13.1-cdh5.3.5
	org.apache.hive	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.5
	org.apache.hive	hive-hcatalog-pig-adapter	0.13.1-cdh5.3.5
	org.apache.hive	hive-hcatalog-server-extensions	0.13.1-cdh5.3.5
	org.apache.hive	hive-hcatalog-streaming	0.13.1-cdh5.3.5
	org.apache.hive	hive-webhcat	0.13.1-cdh5.3.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hive-webhcat-java-client	0.13.1-cdh5.3.5
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.3.5
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.3.5
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.3.5
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.3.5
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-common	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-examples	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-it	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-server	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-shell	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.3.5
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.3.5
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.3.5
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.5
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.3.5
	org.apache.pig	piggybank	0.12.0-cdh5.3.5
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.5
	org.apache.pig	pigunit	0.12.0-cdh5.3.5
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.5

Project	groupId	artifactId	version
	org.apache.flume	flume-tools	1.5.0-cdh5.3.5
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.5
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.5
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.5
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.5
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.5
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.5
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.5
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.5
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.5
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.5
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.5
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.5
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.5
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.5
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-core	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-examples	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.3.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.3.5
	org.apache.oozie	oozie-tools	4.0.0-cdh5.3.5
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.3.5
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.5
	org.apache.mahout	mahout-core	0.9-cdh5.3.5
	org.apache.mahout	mahout-examples	0.9-cdh5.3.5
	org.apache.mahout	mahout-integration	0.9-cdh5.3.5
	org.apache.mahout	mahout-math	0.9-cdh5.3.5
	org.apache.mahout	mahout-math-scala	0.9-cdh5.3.5
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.5

Project	groupId	artifactId	version
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.5
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.5
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.5
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.5
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.5
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.5
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.5
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.5
	com.twitter	parquet-cascading	1.5.0-cdh5.3.5
	com.twitter	parquet-column	1.5.0-cdh5.3.5
	com.twitter	parquet-common	1.5.0-cdh5.3.5
	com.twitter	parquet-encoding	1.5.0-cdh5.3.5
	com.twitter	parquet-format	2.1.0-cdh5.3.5
	com.twitter	parquet-generator	1.5.0-cdh5.3.5
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.5
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.5
	com.twitter	parquet-jackson	1.5.0-cdh5.3.5
	com.twitter	parquet-pig	1.5.0-cdh5.3.5
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.5
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.5
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.3.5
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.3.5
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.5
	com.twitter	parquet-thrift	1.5.0-cdh5.3.5
	com.twitter	parquet-tools	1.5.0-cdh5.3.5
Llama	com.doudera.llama	llama	1.0.0-cdh5.3.5
Apache Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.5

Project	groupId	artifactId	version
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-streaming_2.10	1.2.0-cdh5.3.5
	org.apache.spark	spark-yarn_2.10	1.2.0-cdh5.3.5
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.3.5
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.3.5
	org.apache.crunch	crunch-core	0.11.0-cdh5.3.5
	org.apache.crunch	crunch-examples	0.11.0-cdh5.3.5
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.3.5
	org.apache.crunch	crunch-hive	0.11.0-cdh5.3.5
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.3.5
	org.apache.crunch	crunch-spark	0.11.0-cdh5.3.5
	org.apache.crunch	crunch-test	0.11.0-cdh5.3.5
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.3.5
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.5
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.5
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.5
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.5
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.5
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.5
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	avro-tools	1.7.6-cdh5.3.5
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.5
	org.apache.avro	trevni-core	1.7.6-cdh5.3.5
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.3.5
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.3.5
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.3.5
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.5
	org.kitesdk	kite-data-hive	0.15.0-cdh5.3.5
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.3.5
	org.kitesdk	kite-data-spark	0.15.0-cdh5.3.5
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.3.5
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.3.5
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.3.5
	org.kitesdk	kite-tools	0.15.0-cdh5.3.5

Project	groupId	artifactId	version
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.5
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.5
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.5
	org.apache.solr	solr-cell	4.4.0-cdh5.3.5
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.5
	org.apache.solr	solr-core	4.4.0-cdh5.3.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.5
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.5
	org.apache.solr	solr-langid	4.4.0-cdh5.3.5
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.5
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.5
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.5
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.5
	org.apache.solr	solr-uima	4.4.0-cdh5.3.5
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.5
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.3.5
	com.cloudera.search	search-mr	1.0.0-cdh5.3.5
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.5
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.5
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.5
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.5
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.5
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.5
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.5
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.5
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.5
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.5
	com.ngdata	hbase-sep-api	1.5-cdh5.3.5
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.5
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.5
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.5
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.5

CDH 5.3.4

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.4

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-aws	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-azure	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.3.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.3.4
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.4
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	250-mr1-cdh5.3.4
	org.apache.hadoop	hadoop-core	250-mr1-cdh5.3.4
	org.apache.hadoop	hadoop-examples	250-mr1-cdh5.3.4
	org.apache.hadoop	hadoop-minicluster	250-mr1-cdh5.3.4
	org.apache.hadoop	hadoop-streaming	250-mr1-cdh5.3.4
	org.apache.hadoop	hadoop-test	250-mr1-cdh5.3.4
	org.apache.hadoop	hadoop-tools	250-mr1-cdh5.3.4
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.3.4
	org.apache.hive	hive-beeline	0.13.1-cdh5.3.4
	org.apache.hive	hive-cli	0.13.1-cdh5.3.4
	org.apache.hive	hive-common	0.13.1-cdh5.3.4
	org.apache.hive	hive-contrib	0.13.1-cdh5.3.4
	org.apache.hive	hive-exec	0.13.1-cdh5.3.4
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.3.4
	org.apache.hive	hive-hwi	0.13.1-cdh5.3.4
	org.apache.hive	hive-jdbc	0.13.1-cdh5.3.4
	org.apache.hive	hive-metastore	0.13.1-cdh5.3.4
	org.apache.hive	hive-serde	0.13.1-cdh5.3.4
	org.apache.hive	hive-service	0.13.1-cdh5.3.4
	org.apache.hive	hive-shims	0.13.1-cdh5.3.4
	org.apache.hive	hive-testutils	0.13.1-cdh5.3.4

Project	groupId	artifactId	version
	org.apache.hcatalog	hive-hcatalog-core	0.13.1-cdh5.3.4
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.4
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.13.1-cdh5.3.4
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.3.4
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.3.4
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.3.4
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.3.4
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.3.4
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.3.4
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.3.4
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.3.4
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-common	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-examples	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-it	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-server	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-shell	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.3.4
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.3.4
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.3.4
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.4
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.3.4
	org.apache.pig	piggybank	0.12.0-cdh5.3.4
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.4
	org.apache.pig	pigunit	0.12.0-cdh5.3.4

Version and Download Information

Project	groupId	artifactId	version	
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.4	
	org.apache.flume	flume-tools	1.5.0-cdh5.3.4	
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.4	
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.4	
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.4	
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.4	
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.4	
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.4	
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.4	
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.4	
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.4	
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.4	
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.4	
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.4	
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.4	
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.4	
	Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.4
		org.apache.oozie	oozie-core	4.0.0-cdh5.3.4
org.apache.oozie		oozie-examples	4.0.0-cdh5.3.4	
org.apache.oozie		oozie-hadoop	4.0.0-cdh5.3.4	
org.apache.oozie		oozie-hadoop-distcp	4.0.0-cdh5.3.4	

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-tools	4.0.0-cdh5.3.4
	org.apache.oozie	oozie-mini	4.0.0-cdh5.3.4
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.4
	org.apache.mahout	mahout-core	0.9-cdh5.3.4
	org.apache.mahout	mahout-examples	0.9-cdh5.3.4
	org.apache.mahout	mahout-integration	0.9-cdh5.3.4
	org.apache.mahout	mahout-math	0.9-cdh5.3.4
	org.apache.mahout	mahout-math-scala	0.9-cdh5.3.4
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.4
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.4
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.4
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.4
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.4
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.4

Project	groupId	artifactId	version
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.4
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.4
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.4
	com.twitter	parquet-cascading	1.5.0-cdh5.3.4
	com.twitter	parquet-column	1.5.0-cdh5.3.4
	com.twitter	parquet-common	1.5.0-cdh5.3.4
	com.twitter	parquet-encoding	1.5.0-cdh5.3.4
	com.twitter	parquet-format	2.1.0-cdh5.3.4
	com.twitter	parquet-generator	1.5.0-cdh5.3.4
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.4
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.4
	com.twitter	parquet-jackson	1.5.0-cdh5.3.4
	com.twitter	parquet-pig	1.5.0-cdh5.3.4
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.4
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.4
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.3.4
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.3.4
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.4
	com.twitter	parquet-thrift	1.5.0-cdh5.3.4
	com.twitter	parquet-tools	1.5.0-cdh5.3.4

Version and Download Information

Project	groupId	artifactId	version
Llama	com.doudera.llama	llama	1.0.0-cdh5.3.4
Apache Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.4
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.4
	Apache Crunch	org.apache.spark	spark-streaming_2.10
org.apache.spark		spark-yarn_2.10	1.2.0-cdh5.3.4
org.apache.crunch		crunch-archetype	0.11.0-cdh5.3.4
org.apache.crunch		crunch-contrib	0.11.0-cdh5.3.4
org.apache.crunch		crunch-core	0.11.0-cdh5.3.4
org.apache.crunch		crunch-examples	0.11.0-cdh5.3.4
org.apache.crunch		crunch-hbase	0.11.0-cdh5.3.4
org.apache.crunch		crunch-hive	0.11.0-cdh5.3.4
org.apache.crunch		crunch-scrunch	0.11.0-cdh5.3.4
org.apache.crunch		crunch-spark	0.11.0-cdh5.3.4
Apache Avro	org.apache.crunch	crunch-test	0.11.0-cdh5.3.4
	org.apache.avro	avro	1.7.6-cdh5.3.4
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.4

Project	groupId	artifactId	version
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.4
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.4
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.4
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.4
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.4
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.4
	org.apache.avro	avro-tools	1.7.6-cdh5.3.4
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.4
	org.apache.avro	trevni-core	1.7.6-cdh5.3.4
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.3.4
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.3.4
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.3.4
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.4
	org.kitesdk	kite-data-hive	0.15.0-cdh5.3.4
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.3.4
	org.kitesdk	kite-data-spark	0.15.0-cdh5.3.4
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.3.4
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.3.4

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.3.4
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.3.4
	org.kitesdk	kite-tools	0.15.0-cdh5.3.4
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.4

Project	groupId	artifactId	version
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.4
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.4
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.4
	org.apache.solr	solr-cell	4.4.0-cdh5.3.4
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.4
	org.apache.solr	solr-core	4.4.0-cdh5.3.4
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.4
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.4
	org.apache.solr	solr-langid	4.4.0-cdh5.3.4
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.4
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.4
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.4
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.4
	org.apache.solr	solr-uima	4.4.0-cdh5.3.4
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.4
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.3.4
	com.cloudera.search	search-mr	1.0.0-cdh5.3.4
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.4
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.4
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.4
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.4
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.4
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.4
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.4
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.4
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.4
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.4
	com.ngdata	hbase-sep-api	1.5-cdh5.3.4
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.4
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.4

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.4
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.4

CDH 5.3.3

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-aws	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-azure	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.3.3

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.3.3
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.3
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-mr1-cdh5.3.3
	org.apache.hadoop	hadoop-core	2.5.0-mr1-cdh5.3.3
	org.apache.hadoop	hadoop-examples	2.5.0-mr1-cdh5.3.3
	org.apache.hadoop	hadoop-minicluster	2.5.0-mr1-cdh5.3.3
	org.apache.hadoop	hadoop-streaming	2.5.0-mr1-cdh5.3.3
	org.apache.hadoop	hadoop-test	2.5.0-mr1-cdh5.3.3
	org.apache.hadoop	hadoop-tools	2.5.0-mr1-cdh5.3.3
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.3.3
	org.apache.hive	hive-beeline	0.13.1-cdh5.3.3
	org.apache.hive	hive-cli	0.13.1-cdh5.3.3
	org.apache.hive	hive-common	0.13.1-cdh5.3.3
	org.apache.hive	hive-contrib	0.13.1-cdh5.3.3
	org.apache.hive	hive-exec	0.13.1-cdh5.3.3
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.3.3
	org.apache.hive	hive-hwi	0.13.1-cdh5.3.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-jdbc	0.13.1-cdh5.3.3
	org.apache.hive	hive-metastore	0.13.1-cdh5.3.3
	org.apache.hive	hive-serde	0.13.1-cdh5.3.3
	org.apache.hive	hive-service	0.13.1-cdh5.3.3
	org.apache.hive	hive-shims	0.13.1-cdh5.3.3
	org.apache.hive	hive-testutils	0.13.1-cdh5.3.3
	org.apache.hcatalog	hive-hcatalog-core	0.13.1-cdh5.3.3
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.3
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.13.1-cdh5.3.3
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.3.3
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.3.3
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.3.3
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.3.3
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.3.3
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.3.3
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.3.3
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.3.3
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-common	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-examples	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-it	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-server	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-shell	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.3.3
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.3.3
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.3.3

Project	groupId	artifactId	version
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.3
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.3.3
	org.apache.pig	piggybank	0.12.0-cdh5.3.3
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.3
	org.apache.pig	pigunit	0.12.0-cdh5.3.3
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.3
	org.apache.flume	flume-tools	1.5.0-cdh5.3.3
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.3
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.3
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.3
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.3
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.3
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.3
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.3
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.3
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.3
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.3
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.3
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.3
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.3
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.3

Version and Download Information

Project	groupId	artifactId	version	
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-core	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-examples	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-tools	4.0.0-cdh5.3.3	
	org.apache.oozie	oozie-test	oozie-mini	4.0.0-cdh5.3.3
	Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.3
org.apache.mahout		mahout-core	0.9-cdh5.3.3	
org.apache.mahout		mahout-examples	0.9-cdh5.3.3	
org.apache.mahout		mahout-integration	0.9-cdh5.3.3	
org.apache.mahout		mahout-math	0.9-cdh5.3.3	
org.apache.mahout		mahout-math-scala	0.9-cdh5.3.3	
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.3	
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.3	
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.3	
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.3	
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.3	

Project	groupId	artifactId	version
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.3
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.3
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.3
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.3
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.3
Apache Sentry	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.3
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.3
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.3
	com.twitter	parquet-cascading	1.5.0-cdh5.3.3
	com.twitter	parquet-column	1.5.0-cdh5.3.3
	com.twitter	parquet-common	1.5.0-cdh5.3.3
	com.twitter	parquet-encoding	1.5.0-cdh5.3.3
	com.twitter	parquet-format	2.1.0-cdh5.3.3
	com.twitter	parquet-generator	1.5.0-cdh5.3.3
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.3
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.3
	com.twitter	parquet-jackson	1.5.0-cdh5.3.3
	com.twitter	parquet-pig	1.5.0-cdh5.3.3
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.3
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.3

Project	groupId	artifactId	version
	com.twitter	parquet-scala_2.10	1.5.0-cdh5.3.3
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.3.3
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.3
	com.twitter	parquet-thrift	1.5.0-cdh5.3.3
	com.twitter	parquet-tools	1.5.0-cdh5.3.3
Llama	com.doudera.llama	llama	1.0.0-cdh5.3.3
Apache Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-streaming_2.10	1.2.0-cdh5.3.3
	org.apache.spark	spark-yarn_2.10	1.2.0-cdh5.3.3
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.3.3
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.3.3
	org.apache.crunch	crunch-core	0.11.0-cdh5.3.3
	org.apache.crunch	crunch-examples	0.11.0-cdh5.3.3
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.3.3
	org.apache.crunch	crunch-hive	0.11.0-cdh5.3.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.3.3
	org.apache.crunch	crunch-spark	0.11.0-cdh5.3.3
	org.apache.crunch	crunch-test	0.11.0-cdh5.3.3
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.3.3
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.3
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.3
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.3
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.3
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.3
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.3
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.3
	org.apache.avro	avro-tools	1.7.6-cdh5.3.3
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.3
	org.apache.avro	trevni-core	1.7.6-cdh5.3.3
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.3.3
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.3.3
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.3.3
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.3
	org.kitesdk	kite-data-hive	0.15.0-cdh5.3.3
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.3.3
	org.kitesdk	kite-data-spark	0.15.0-cdh5.3.3
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.3.3
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.3.3

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.3.3
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.3.3
	org.kitesdk	kite-tools	0.15.0-cdh5.3.3
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.3
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.3
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.3
	org.apache.solr	solr-cell	4.4.0-cdh5.3.3
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.3
	org.apache.solr	solr-core	4.4.0-cdh5.3.3
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.3
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.3
	org.apache.solr	solr-langid	4.4.0-cdh5.3.3
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.3
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.3
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.3
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.3
	org.apache.solr	solr-uima	4.4.0-cdh5.3.3
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.3
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.3.3
	com.cloudera.search	search-mr	1.0.0-cdh5.3.3
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.3
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.3
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.3
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.3
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.3
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.3
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.3
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.3

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.3
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.3
	com.ngdata	hbase-sep-api	1.5-cdh5.3.3
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.3
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.3
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.3
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.3

CDH 5.3.2

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-aws	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-azure	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.3.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.3.2
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.2
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-mr1-cdh5.3.2
	org.apache.hadoop	hadoop-core	2.5.0-mr1-cdh5.3.2
	org.apache.hadoop	hadoop-examples	2.5.0-mr1-cdh5.3.2
	org.apache.hadoop	hadoop-minicluster	2.5.0-mr1-cdh5.3.2
	org.apache.hadoop	hadoop-streaming	2.5.0-mr1-cdh5.3.2
	org.apache.hadoop	hadoop-test	2.5.0-mr1-cdh5.3.2
	org.apache.hadoop	hadoop-tools	2.5.0-mr1-cdh5.3.2
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.3.2
	org.apache.hive	hive-beeline	0.13.1-cdh5.3.2
	org.apache.hive	hive-cli	0.13.1-cdh5.3.2

Project	groupId	artifactId	version
	org.apache.hive	hive-common	0.13.1-cdh5.3.2
	org.apache.hive	hive-contrib	0.13.1-cdh5.3.2
	org.apache.hive	hive-exec	0.13.1-cdh5.3.2
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.3.2
	org.apache.hive	hive-hwi	0.13.1-cdh5.3.2
	org.apache.hive	hive-jdbc	0.13.1-cdh5.3.2
	org.apache.hive	hive-metastore	0.13.1-cdh5.3.2
	org.apache.hive	hive-serde	0.13.1-cdh5.3.2
	org.apache.hive	hive-service	0.13.1-cdh5.3.2
	org.apache.hive	hive-shims	0.13.1-cdh5.3.2
	org.apache.hive	hive-testutils	0.13.1-cdh5.3.2
	org.apache.hcatalog	hive-hcatalog-core	0.13.1-cdh5.3.2
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.2
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.13.1-cdh5.3.2
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.3.2
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.3.2
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.3.2
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.3.2
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.3.2
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.3.2
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.3.2
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.3.2
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-common	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-examples	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-it	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.3.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-server	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-shell	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.3.2
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.3.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.3.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.3.2
	org.apache.pig	piggybank	0.12.0-cdh5.3.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.2
	org.apache.pig	pigunit	0.12.0-cdh5.3.2
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.2
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.2
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.2
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.2
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.2
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.2
	org.apache.flume	flume-tools	1.5.0-cdh5.3.2
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.2
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.2
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.2
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.2
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.2
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.2
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.2
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.2
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.2
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.2
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.2
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.2

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.2
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.2
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.2
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.2
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.2
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.2
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-core	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-examples	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-tools	4.0.0-cdh5.3.2
	org.apache.oozie	oozie-mini	4.0.0-cdh5.3.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.2
	org.apache.mahout	mahout-core	0.9-cdh5.3.2
	org.apache.mahout	mahout-examples	0.9-cdh5.3.2
	org.apache.mahout	mahout-integration	0.9-cdh5.3.2
	org.apache.mahout	mahout-math	0.9-cdh5.3.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.mahout	mahout-math-scala	0.9-cdh5.3.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.2
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.2

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.2
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.2
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.2
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.2
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.2
	com.twitter	parquet-cascading	1.5.0-cdh5.3.2
	com.twitter	parquet-column	1.5.0-cdh5.3.2
	com.twitter	parquet-common	1.5.0-cdh5.3.2
	com.twitter	parquet-encoding	1.5.0-cdh5.3.2
	com.twitter	parquet-format	2.1.0-cdh5.3.2
	com.twitter	parquet-generator	1.5.0-cdh5.3.2

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.2
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.2
	com.twitter	parquet-jackson	1.5.0-cdh5.3.2
	com.twitter	parquet-pig	1.5.0-cdh5.3.2
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.2
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.2
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.3.2
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.2
	com.twitter	parquet-thrift	1.5.0-cdh5.3.2
	com.twitter	parquet-tools	1.5.0-cdh5.3.2
Llama	com.douderallama	llama	1.0.0-cdh5.3.2
Apache Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.2
	org.apache.spark	spark-streaming_2.10	1.2.0-cdh5.3.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.3.2
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.3.2

Project	groupId	artifactId	version
	org.apache.crunch	crunch-core	0.11.0-cdh5.3.2
	org.apache.crunch	crunch-examples	0.11.0-cdh5.3.2
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.3.2
	org.apache.crunch	crunch-hive	0.11.0-cdh5.3.2
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.3.2
	org.apache.crunch	crunch-spark	0.11.0-cdh5.3.2
	org.apache.crunch	crunch-test	0.11.0-cdh5.3.2
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.3.2
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.2
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.2
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.2
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.2
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.2
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.2
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.2
	org.apache.avro	avro-tools	1.7.6-cdh5.3.2
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.2
	org.apache.avro	trevni-core	1.7.6-cdh5.3.2
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.3.2
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.3.2
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.3.2
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.2
	org.kitesdk	kite-data-hive	0.15.0-cdh5.3.2
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.3.2
	org.kitesdk	kite-data-spark	0.15.0-cdh5.3.2
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.3.2
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.3.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.3.2
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.3.2
	org.kitesdk	kite-tools	0.15.0-cdh5.3.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.2
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.2
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.2
	org.apache.solr	solr-cell	4.4.0-cdh5.3.2
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.2
	org.apache.solr	solr-core	4.4.0-cdh5.3.2
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.2
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.2
	org.apache.solr	solr-langid	4.4.0-cdh5.3.2
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.2
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.2
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.2
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.2
	org.apache.solr	solr-uima	4.4.0-cdh5.3.2
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.2
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.3.2
	com.cloudera.search	search-mr	1.0.0-cdh5.3.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.2

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.2
	com.ngdata	hbase-sep-api	1.5-cdh5.3.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.2
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.2

CDH 5.3.1

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-aws	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-azure	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.3.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-mapreduce-client-native-task	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.3.1
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-mr1-cdh5.3.1
	org.apache.hadoop	hadoop-core	2.5.0-mr1-cdh5.3.1
	org.apache.hadoop	hadoop-examples	2.5.0-mr1-cdh5.3.1
	org.apache.hadoop	hadoop-minicluster	2.5.0-mr1-cdh5.3.1
	org.apache.hadoop	hadoop-streaming	2.5.0-mr1-cdh5.3.1
	org.apache.hadoop	hadoop-test	2.5.0-mr1-cdh5.3.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-tools	2.5.0-mr1-cdh5.3.1
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.3.1
	org.apache.hive	hive-beeline	0.13.1-cdh5.3.1
	org.apache.hive	hive-cli	0.13.1-cdh5.3.1
	org.apache.hive	hive-common	0.13.1-cdh5.3.1
	org.apache.hive	hive-contrib	0.13.1-cdh5.3.1
	org.apache.hive	hive-exec	0.13.1-cdh5.3.1
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.3.1
	org.apache.hive	hive-hwi	0.13.1-cdh5.3.1
	org.apache.hive	hive-jdbc	0.13.1-cdh5.3.1
	org.apache.hive	hive-metastore	0.13.1-cdh5.3.1
	org.apache.hive	hive-serde	0.13.1-cdh5.3.1
	org.apache.hive	hive-service	0.13.1-cdh5.3.1
	org.apache.hive	hive-shims	0.13.1-cdh5.3.1
	org.apache.hive	hive-testutils	0.13.1-cdh5.3.1
	org.apache.hcatalog	hive-hcatalog-core	0.13.1-cdh5.3.1
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.1
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.13.1-cdh5.3.1
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.3.1
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.3.1
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.3.1
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.3.1
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.3.1
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.3.1
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.3.1
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.3.1
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-common	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-examples	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.3.1

Project	groupId	artifactId	version
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-it	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-server	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-shell	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.3.1
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.3.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.3.1
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.3.1
	org.apache.pig	piggybank	0.12.0-cdh5.3.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.1
	org.apache.pig	pigunit	0.12.0-cdh5.3.1
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.1
	org.apache.flume	flume-tools	1.5.0-cdh5.3.1
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.1
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.1
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.1
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.1
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.1
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.1
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.1
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.1
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.1
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.1
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.1
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.1
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.1
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-core	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-examples	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-tools	4.0.0-cdh5.3.1
	org.apache.oozie	oozie-mini	4.0.0-cdh5.3.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.1

Project	groupId	artifactId	version
	org.apache.mahout	mahout-core	0.9-cdh5.3.1
	org.apache.mahout	mahout-examples	0.9-cdh5.3.1
	org.apache.mahout	mahout-integration	0.9-cdh5.3.1
	org.apache.mahout	mahout-math	0.9-cdh5.3.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.3.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.1
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.1
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.1
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.1
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.1
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.1
	com.twitter	parquet-cascading	1.5.0-cdh5.3.1
	com.twitter	parquet-column	1.5.0-cdh5.3.1

Project	groupId	artifactId	version
	com.twitter	parquet-common	1.5.0-cdh5.3.1
	com.twitter	parquet-encoding	1.5.0-cdh5.3.1
	com.twitter	parquet-format	2.1.0-cdh5.3.1
	com.twitter	parquet-generator	1.5.0-cdh5.3.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.1
	com.twitter	parquet-jackson	1.5.0-cdh5.3.1
	com.twitter	parquet-pig	1.5.0-cdh5.3.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.1
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.3.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.1
	com.twitter	parquet-thrift	1.5.0-cdh5.3.1
	com.twitter	parquet-tools	1.5.0-cdh5.3.1
Llama	com.douderallama	llama	1.0.0-cdh5.3.1
Apache Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.1
	org.apache.spark	spark-streaming_2.10	1.2.0-cdh5.3.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.3.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.3.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.3.1
	org.apache.crunch	crunch-examples	0.11.0-cdh5.3.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.3.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.3.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.3.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.3.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.3.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.3.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.1
	org.apache.avro	avro-tools	1.7.6-cdh5.3.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.1
	org.apache.avro	trevni-core	1.7.6-cdh5.3.1
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.3.1
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.3.1
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.3.1
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.1
	org.kitesdk	kite-data-hive	0.15.0-cdh5.3.1
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.3.1
	org.kitesdk	kite-data-spark	0.15.0-cdh5.3.1
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.3.1

Project	groupId	artifactId	version
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.3.1
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.3.1
	org.kitesdk	kite-tools	0.15.0-cdh5.3.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.1
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.1
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.1
	org.apache.solr	solr-cell	4.4.0-cdh5.3.1
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.1
	org.apache.solr	solr-core	4.4.0-cdh5.3.1
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.1
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.1
	org.apache.solr	solr-langid	4.4.0-cdh5.3.1
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.1
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.1
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.1
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.1
	org.apache.solr	solr-uima	4.4.0-cdh5.3.1
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.3.1
	com.cloudera.search	search-mr	1.0.0-cdh5.3.1

Project	groupId	artifactId	version
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.1
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.1
	com.ngdata	hbase-sep-api	1.5-cdh5.3.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.1
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.1

CDH 5.3.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-aws	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-azure	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.3.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.3.0
Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-cdh5.3.0
	org.apache.hadoop	hadoop-core	2.5.0-cdh5.3.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-examples	2.5.0-mr1-cdh5.3.0
	org.apache.hadoop	hadoop-minicluster	2.5.0-mr1-cdh5.3.0
	org.apache.hadoop	hadoop-streaming	2.5.0-mr1-cdh5.3.0
	org.apache.hadoop	hadoop-test	2.5.0-mr1-cdh5.3.0
	org.apache.hadoop	hadoop-tools	2.5.0-mr1-cdh5.3.0
Hive	org.apache.hive	hive-ant	0.13.1-cdh5.3.0
	org.apache.hive	hive-beeline	0.13.1-cdh5.3.0
	org.apache.hive	hive-cli	0.13.1-cdh5.3.0
	org.apache.hive	hive-common	0.13.1-cdh5.3.0
	org.apache.hive	hive-contrib	0.13.1-cdh5.3.0
	org.apache.hive	hive-exec	0.13.1-cdh5.3.0
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.3.0
	org.apache.hive	hive-hwi	0.13.1-cdh5.3.0
	org.apache.hive	hive-jdbc	0.13.1-cdh5.3.0
	org.apache.hive	hive-metastore	0.13.1-cdh5.3.0
	org.apache.hive	hive-serde	0.13.1-cdh5.3.0
	org.apache.hive	hive-service	0.13.1-cdh5.3.0
	org.apache.hive	hive-shims	0.13.1-cdh5.3.0
	org.apache.hive	hive-testutils	0.13.1-cdh5.3.0
	org.apache.hive	hive-hcatalog-core	0.13.1-cdh5.3.0
	org.apache.hive	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.3.0
	org.apache.hive	hive-hcatalog-pig-adapter	0.13.1-cdh5.3.0
	org.apache.hive	hive-hcatalog-server-extensions	0.13.1-cdh5.3.0
	org.apache.hive	hive-hcatalog-streaming	0.13.1-cdh5.3.0
	org.apache.hive	hive-webhcat	0.13.1-cdh5.3.0
	org.apache.hive	hive-webhcat-java-client	0.13.1-cdh5.3.0
	org.apache.hive	hive-shims-0.23	0.13.1-cdh5.3.0
	org.apache.hive	hive-shims-common	0.13.1-cdh5.3.0
	org.apache.hive	hive-shims-common-secure	0.13.1-cdh5.3.0
	org.apache.hive	hive-shims-scheduler	0.13.1-cdh5.3.0

Version and Download Information

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-common	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-examples	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-it	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-server	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-shell	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.3.0
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.3.0
	ZooKeeper	org.apache.zookeeper	zookeeper
Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.3.0
Pig	org.apache.pig	pig	0.12.0-cdh5.3.0
	org.apache.pig	piggybank	0.12.0-cdh5.3.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.3.0
	org.apache.pig	pigunit	0.12.0-cdh5.3.0
Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-core	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-node	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.3.0
	org.apache.flume	flume-tools	1.5.0-cdh5.3.0
	org.apache.flume	flume-file-channel	1.5.0-cdh5.3.0
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.3.0
	org.apache.flume	flume-kafka-channel	1.5.0-cdh5.3.0
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.3.0

Project	groupId	artifactId	version
	org.apache.flume	flume-avro-source	1.5.0-cdh5.3.0
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.3.0
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.3.0
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.3.0
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.3.0
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.3.0
	org.apache.flume	flume-jms-source	1.5.0-cdh5.3.0
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.3.0
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.3.0
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.3.0
Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-core	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-examples	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.3.0
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.3.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-tools	4.0.0-cdh5.3.0
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.3.0
Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.3.0
	org.apache.mahout	mahout-core	0.9-cdh5.3.0
	org.apache.mahout	mahout-examples	0.9-cdh5.3.0
	org.apache.mahout	mahout-integration	0.9-cdh5.3.0
	org.apache.mahout	mahout-math	0.9-cdh5.3.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.3.0
Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.3.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.3.0
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.3.0
Sqoop2	org.apache.sqoop	connector-sdk	1.99.4-cdh5.3.0

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-client	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-common	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-core	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-docs	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-security	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-shell	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-spi	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-tomcat	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-tools	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-connector-hdfs	1.99.4-cdh5.3.0
	org.apache.sqoop	sqoop-repository-derby	1.99.4-cdh5.3.0
Sentry	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-dist	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-hdfs-common	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-hdfs-dist	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-hdfs-namenode-plugin	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.3.0
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.3.0

Version and Download Information

Project	groupId	artifactId	version
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.3.0
	com.twitter	parquet-cascading	1.5.0-cdh5.3.0
	com.twitter	parquet-column	1.5.0-cdh5.3.0
	com.twitter	parquet-common	1.5.0-cdh5.3.0
	com.twitter	parquet-encoding	1.5.0-cdh5.3.0
	com.twitter	parquet-format	2.1.0-cdh5.3.0
	com.twitter	parquet-generator	1.5.0-cdh5.3.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.3.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.3.0
	com.twitter	parquet-jackson	1.5.0-cdh5.3.0
	com.twitter	parquet-pig	1.5.0-cdh5.3.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.3.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.3.0
	com.twitter	parquet-scrooge_2.10	1.5.0-cdh5.3.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.3.0
	com.twitter	parquet-thrift	1.5.0-cdh5.3.0
	com.twitter	parquet-tools	1.5.0-cdh5.3.0
Llama	com.doudera.llama	llama	1.0.0-cdh5.3.0
Spark	org.apache.spark	spark-bagel_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-catalyst_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-core_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-graphx_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-hive_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-mllib_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-network-common_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-network-shuffle_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-network-yarn_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-sql_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-streaming-flume_2.10	1.2.0-cdh5.3.0

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-kafka_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-streaming-twitter_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.2.0-cdh5.3.0
	org.apache.spark	spark-streaming_2.10	1.2.0-cdh5.3.0
Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.3.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.3.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.3.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.3.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.3.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.3.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.3.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.3.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.3.0
Avro	org.apache.avro	avro	1.7.6-cdh5.3.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.3.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.3.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.3.0
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.3.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.3.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.3.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.3.0
	org.apache.avro	avro-tools	1.7.6-cdh5.3.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.3.0
	org.apache.avro	trevni-core	1.7.6-cdh5.3.0
Kite	org.kitesdk	kite-data-core	0.15.0-cdh5.3.0
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.3.0
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.3.0
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.3.0
	org.kitesdk	kite-data-hive	0.15.0-cdh5.3.0

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.3.0
	org.kitesdk	kite-data-spark	0.15.0-cdh5.3.0
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.3.0
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.3.0
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.3.0
	org.kitesdk	kite-tools	0.15.0-cdh5.3.0
Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.3.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-classification	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-core	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-demo	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-facet	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-join	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-memory	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-misc	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-queries	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.3.0
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.3.0
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.3.0
	org.apache.solr	solr-cell	4.4.0-cdh5.3.0
	org.apache.solr	solr-clustering	4.4.0-cdh5.3.0
	org.apache.solr	solr-core	4.4.0-cdh5.3.0
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.3.0
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.3.0
	org.apache.solr	solr-langid	4.4.0-cdh5.3.0
	org.apache.solr	solr-security-util	4.4.0-cdh5.3.0
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.3.0
	org.apache.solr	solr-solrj	4.4.0-cdh5.3.0
	org.apache.solr	solr-test-framework	4.4.0-cdh5.3.0
	org.apache.solr	solr-uima	4.4.0-cdh5.3.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-velocity	4.4.0-cdh5.3.0
Search	com.cloudera.search	search-crunch	1.0.0-cdh5.3.0
	com.cloudera.search	search-mr	1.0.0-cdh5.3.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.3.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.3.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.3.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.3.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.3.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.3.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.3.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.3.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.3.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.3.0
	com.ngdata	hbase-sep-api	1.5-cdh5.3.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.3.0
	com.ngdata	hbase-sep-impl	1.5-cdh5.3.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.3.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.3.0

Maven Artifacts for CDH 5.2.x Releases

CDH 5.2.6

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.2.6

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-mapreduce-client-native-task	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.2.6
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.2.6
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-mr1-cdh5.2.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-core	2.5.0-mr1-cdh5.2.6
	org.apache.hadoop	hadoop-examples	2.5.0-mr1-cdh5.2.6
	org.apache.hadoop	hadoop-minicluster	2.5.0-mr1-cdh5.2.6
	org.apache.hadoop	hadoop-streaming	2.5.0-mr1-cdh5.2.6
	org.apache.hadoop	hadoop-test	2.5.0-mr1-cdh5.2.6
	org.apache.hadoop	hadoop-tools	2.5.0-mr1-cdh5.2.6
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.2.6
	org.apache.hive	hive-beeline	0.13.1-cdh5.2.6
	org.apache.hive	hive-cli	0.13.1-cdh5.2.6
	org.apache.hive	hive-common	0.13.1-cdh5.2.6
	org.apache.hive	hive-contrib	0.13.1-cdh5.2.6
	org.apache.hive	hive-exec	0.13.1-cdh5.2.6
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.2.6
	org.apache.hive	hive-hwi	0.13.1-cdh5.2.6
	org.apache.hive	hive-jdbc	0.13.1-cdh5.2.6
	org.apache.hive	hive-metastore	0.13.1-cdh5.2.6
	org.apache.hive	hive-serde	0.13.1-cdh5.2.6
	org.apache.hive	hive-service	0.13.1-cdh5.2.6
	org.apache.hive	hive-shims	0.13.1-cdh5.2.6
	org.apache.hive	hive-testutils	0.13.1-cdh5.2.6
	org.apache.hcatalog	hive-hcatalog-core	0.13.1-cdh5.2.6
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.2.6
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.13.1-cdh5.2.6
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.2.6
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.2.6
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.2.6
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.2.6
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.2.6
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.2.6
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.2.6

Project	groupId	artifactId	version
	org.apache.hive.shims	hive-shims-scheduler	0.13.1-cdh5.2.6
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-common	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-examples	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-it	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-server	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-shell	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.2.6
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.2.6
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.2.6
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.2.6
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.2.6
	org.apache.pig	piggybank	0.12.0-cdh5.2.6
	org.apache.pig	pigsmoke	0.12.0-cdh5.2.6
	org.apache.pig	pigunit	0.12.0-cdh5.2.6
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.2.6
	org.apache.flume	flume-ng-core	1.5.0-cdh5.2.6
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.2.6
	org.apache.flume	flume-ng-node	1.5.0-cdh5.2.6
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.2.6
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.2.6
	org.apache.flume	flume-tools	1.5.0-cdh5.2.6
	org.apache.flume	flume-file-channel	1.5.0-cdh5.2.6
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.2.6
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.2.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.2.6
	org.apache.flume	flume-avro-source	1.5.0-cdh5.2.6
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.2.6
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.2.6
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.2.6
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.2.6
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.2.6
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.2.6
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.2.6
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.2.6
	org.apache.flume	flume-jms-source	1.5.0-cdh5.2.6
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.2.6
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.2.6
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.2.6
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-core	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-examples	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.2.6

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.2.6
	org.apache.oozie	oozie-tools	4.0.0-cdh5.2.6
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.2.6
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.2.6
	org.apache.mahout	mahout-core	0.9-cdh5.2.6
	org.apache.mahout	mahout-examples	0.9-cdh5.2.6
	org.apache.mahout	mahout-integration	0.9-cdh5.2.6
	org.apache.mahout	mahout-math	0.9-cdh5.2.6
	org.apache.mahout	mahout-math-scala	0.9-cdh5.2.6
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-chef	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-cli	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-core	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-druid	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-examples	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-hama	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-pig	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-solr	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.2.6
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.2.6
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.2.6

Version and Download Information

Project	groupId	artifactId	version	
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.2.6	
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.2.6	
	Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.2.6
		org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.2.6
org.apache.sentry		sentry-core-common	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-core-model-db	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-core-model-search	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-dist	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-policy-common	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-policy-db	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-policy-search	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-provider-cache	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-provider-common	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-provider-db	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-provider-file	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-tests-hive	1.4.0-cdh5.2.6	
org.apache.sentry		sentry-tests-solr	1.4.0-cdh5.2.6	
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.2.6	
	com.twitter	parquet-cascading	1.5.0-cdh5.2.6	
	com.twitter	parquet-column	1.5.0-cdh5.2.6	
	com.twitter	parquet-common	1.5.0-cdh5.2.6	

Project	groupId	artifactId	version
	com.twitter	parquet-encoding	1.5.0-cdh5.2.6
	com.twitter	parquet-format	2.1.0-cdh5.2.6
	com.twitter	parquet-generator	1.5.0-cdh5.2.6
	com.twitter	parquet-hadoop	1.5.0-cdh5.2.6
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.2.6
	com.twitter	parquet-jackson	1.5.0-cdh5.2.6
	com.twitter	parquet-pig	1.5.0-cdh5.2.6
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.2.6
	com.twitter	parquet-protobuf	1.5.0-cdh5.2.6
	com.twitter	parquet-scrooge	1.5.0-cdh5.2.6
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.2.6
	com.twitter	parquet-thrift	1.5.0-cdh5.2.6
	com.twitter	parquet-tools	1.5.0-cdh5.2.6
Llama	com.doudera.llama	llama	1.0.0-cdh5.2.6
Apache Spark	org.apache.spark	spark-bagel_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-catalyst_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-core_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-graphx_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-hive_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-mllib_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-sql_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-streaming-flume-sink_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-streaming-flume_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-streaming-kafka_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-streaming-mqtt_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-streaming-twitter_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-streaming-zeromq_2.10	1.1.0-cdh5.2.6
	org.apache.spark	spark-streaming_2.10	1.1.0-cdh5.2.6
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.2.6
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.2.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.crunch	crunch-core	0.11.0-cdh5.2.6
	org.apache.crunch	crunch-examples	0.11.0-cdh5.2.6
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.2.6
	org.apache.crunch	crunch-hive	0.11.0-cdh5.2.6
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.2.6
	org.apache.crunch	crunch-spark	0.11.0-cdh5.2.6
	org.apache.crunch	crunch-test	0.11.0-cdh5.2.6
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.2.6
	org.apache.avro	avro-compiler	1.7.6-cdh5.2.6
	org.apache.avro	avro-ipc	1.7.6-cdh5.2.6
	org.apache.avro	avro-mapred	1.7.6-cdh5.2.6
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.2.6
	org.apache.avro	avro-protobuf	1.7.6-cdh5.2.6
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.2.6
	org.apache.avro	avro-thrift	1.7.6-cdh5.2.6
	org.apache.avro	avro-tools	1.7.6-cdh5.2.6
	org.apache.avro	trevni-avro	1.7.6-cdh5.2.6
	org.apache.avro	trevni-core	1.7.6-cdh5.2.6
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.2.6
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.2.6
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.2.6
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.2.6
	org.kitesdk	kite-data-hive	0.15.0-cdh5.2.6
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.2.6
	org.kitesdk	kite-data-spark	0.15.0-cdh5.2.6
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.2.6
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.2.6

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.2.6
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.2.6
	org.kitesdk	kite-tools	0.15.0-cdh5.2.6
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-classification	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-core	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-demo	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-facet	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.2.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-join	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-memory	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-misc	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-queries	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.2.6
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.2.6
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.2.6
	org.apache.solr	solr-cell	4.4.0-cdh5.2.6
	org.apache.solr	solr-clustering	4.4.0-cdh5.2.6
	org.apache.solr	solr-core	4.4.0-cdh5.2.6
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.2.6
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.2.6
	org.apache.solr	solr-langid	4.4.0-cdh5.2.6
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.2.6
	org.apache.solr	solr-solrj	4.4.0-cdh5.2.6
	org.apache.solr	solr-test-framework	4.4.0-cdh5.2.6
	org.apache.solr	solr-uima	4.4.0-cdh5.2.6
	org.apache.solr	solr-velocity	4.4.0-cdh5.2.6
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.2.6
	com.cloudera.search	search-mr	1.0.0-cdh5.2.6
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.2.6
	com.ngdata	hbase-indexer-cli	1.5-cdh5.2.6
	com.ngdata	hbase-indexer-common	1.5-cdh5.2.6
	com.ngdata	hbase-indexer-demo	1.5-cdh5.2.6
	com.ngdata	hbase-indexer-dist	1.5-cdh5.2.6

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-engine	1.5-cdh5.2.6
	com.ngdata	hbase-indexer-model	1.5-cdh5.2.6
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.2.6
	com.ngdata	hbase-indexer-mr	1.5-cdh5.2.6
	com.ngdata	hbase-indexer-server	1.5-cdh5.2.6
	com.ngdata	hbase-sep-api	1.5-cdh5.2.6
	com.ngdata	hbase-sep-demo	1.5-cdh5.2.6
	com.ngdata	hbase-sep-impl	1.5-cdh5.2.6
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.2.6
	com.ngdata	hbase-sep-tools	1.5-cdh5.2.6

CDH 5.2.5

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.2.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.2.5
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.2.5
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-mr1-cdh5.2.5
	org.apache.hadoop	hadoop-core	2.5.0-mr1-cdh5.2.5
	org.apache.hadoop	hadoop-examples	2.5.0-mr1-cdh5.2.5
	org.apache.hadoop	hadoop-minicluster	2.5.0-mr1-cdh5.2.5
	org.apache.hadoop	hadoop-streaming	2.5.0-mr1-cdh5.2.5
	org.apache.hadoop	hadoop-test	2.5.0-mr1-cdh5.2.5
	org.apache.hadoop	hadoop-tools	2.5.0-mr1-cdh5.2.5
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.2.5
	org.apache.hive	hive-beeline	0.13.1-cdh5.2.5

Project	groupId	artifactId	version
	org.apache.hive	hive-cli	0.13.1-cdh5.2.5
	org.apache.hive	hive-common	0.13.1-cdh5.2.5
	org.apache.hive	hive-contrib	0.13.1-cdh5.2.5
	org.apache.hive	hive-exec	0.13.1-cdh5.2.5
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.2.5
	org.apache.hive	hive-hwi	0.13.1-cdh5.2.5
	org.apache.hive	hive-jdbc	0.13.1-cdh5.2.5
	org.apache.hive	hive-metastore	0.13.1-cdh5.2.5
	org.apache.hive	hive-serde	0.13.1-cdh5.2.5
	org.apache.hive	hive-service	0.13.1-cdh5.2.5
	org.apache.hive	hive-shims	0.13.1-cdh5.2.5
	org.apache.hive	hive-testutils	0.13.1-cdh5.2.5
	org.apache.hcatalog	hive-hcatalog-core	0.13.1-cdh5.2.5
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.2.5
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.13.1-cdh5.2.5
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.2.5
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.2.5
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.2.5
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.2.5
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.2.5
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.2.5
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.2.5
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.2.5
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-common	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-examples	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-it	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.2.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-server	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-shell	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.2.5
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.2.5
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.2.5
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.2.5
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.2.5
	org.apache.pig	piggybank	0.12.0-cdh5.2.5
	org.apache.pig	pigsmoke	0.12.0-cdh5.2.5
	org.apache.pig	pigunit	0.12.0-cdh5.2.5
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.2.5
	org.apache.flume	flume-ng-core	1.5.0-cdh5.2.5
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.2.5
	org.apache.flume	flume-ng-node	1.5.0-cdh5.2.5
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.2.5
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.2.5
	org.apache.flume	flume-tools	1.5.0-cdh5.2.5
	org.apache.flume	flume-file-channel	1.5.0-cdh5.2.5
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.2.5
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.2.5
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.2.5
	org.apache.flume	flume-avro-source	1.5.0-cdh5.2.5
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.2.5
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.2.5
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.2.5
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.2.5
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.2.5
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.2.5

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.2.5
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.2.5
	org.apache.flume	flume-jms-source	1.5.0-cdh5.2.5
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.2.5
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.2.5
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.2.5
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-core	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-examples	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-tools	4.0.0-cdh5.2.5
	org.apache.oozie	oozie-mini	4.0.0-cdh5.2.5
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.2.5
	org.apache.mahout	mahout-core	0.9-cdh5.2.5
	org.apache.mahout	mahout-examples	0.9-cdh5.2.5
	org.apache.mahout	mahout-integration	0.9-cdh5.2.5
	org.apache.mahout	mahout-math	0.9-cdh5.2.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.mahout	mahout-math-scala	0.9-cdh5.2.5
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-chef	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-cli	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-core	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-druid	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-examples	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-hama	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-pig	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-solr	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.2.5
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.2.5
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.2.5
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.2.5
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.2.5
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.2.5
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.2.5
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.2.5
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.2.5
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.2.5
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.2.5

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.2.5
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.2.5
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.2.5
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-dist	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.2.5
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.2.5
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.2.5
	com.twitter	parquet-cascading	1.5.0-cdh5.2.5
	com.twitter	parquet-column	1.5.0-cdh5.2.5
	com.twitter	parquet-common	1.5.0-cdh5.2.5
	com.twitter	parquet-encoding	1.5.0-cdh5.2.5
	com.twitter	parquet-format	2.1.0-cdh5.2.5
	com.twitter	parquet-generator	1.5.0-cdh5.2.5
	com.twitter	parquet-hadoop	1.5.0-cdh5.2.5
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.2.5
	com.twitter	parquet-jackson	1.5.0-cdh5.2.5
	com.twitter	parquet-pig	1.5.0-cdh5.2.5
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.2.5

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-protobuf	1.5.0-cdh5.2.5
	com.twitter	parquet-scrooge	1.5.0-cdh5.2.5
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.2.5
	com.twitter	parquet-thrift	1.5.0-cdh5.2.5
	com.twitter	parquet-tools	1.5.0-cdh5.2.5
Llama	com.doudera.llama	llama	1.0.0-cdh5.2.5
Apache Spark	org.apache.spark	spark-bagel_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-catalyst_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-core_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-graphx_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-hive_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-mllib_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-sql_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-streaming-flume-sink_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-streaming-flume_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-streaming-kafka_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-streaming-mqtt_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-streaming-twitter_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-streaming-zeromq_2.10	1.1.0-cdh5.2.5
	org.apache.spark	spark-streaming_2.10	1.1.0-cdh5.2.5
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.2.5
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.2.5
	org.apache.crunch	crunch-core	0.11.0-cdh5.2.5
	org.apache.crunch	crunch-examples	0.11.0-cdh5.2.5
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.2.5
	org.apache.crunch	crunch-hive	0.11.0-cdh5.2.5
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.2.5
	org.apache.crunch	crunch-spark	0.11.0-cdh5.2.5
	org.apache.crunch	crunch-test	0.11.0-cdh5.2.5
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.2.5

Project	groupId	artifactId	version
	org.apache.avro	avro-compiler	1.7.6-cdh5.2.5
	org.apache.avro	avro-ipc	1.7.6-cdh5.2.5
	org.apache.avro	avro-mapred	1.7.6-cdh5.2.5
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.2.5
	org.apache.avro	avro-protobuf	1.7.6-cdh5.2.5
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.2.5
	org.apache.avro	avro-thrift	1.7.6-cdh5.2.5
	org.apache.avro	avro-tools	1.7.6-cdh5.2.5
	org.apache.avro	trevni-avro	1.7.6-cdh5.2.5
	org.apache.avro	trevni-core	1.7.6-cdh5.2.5
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.2.5
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.2.5
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.2.5
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.2.5
	org.kitesdk	kite-data-hive	0.15.0-cdh5.2.5
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.2.5
	org.kitesdk	kite-data-spark	0.15.0-cdh5.2.5
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.2.5
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.2.5

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.2.5
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.2.5
	org.kitesdk	kite-tools	0.15.0-cdh5.2.5
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-classification	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-core	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-demo	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-facet	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-join	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-memory	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-misc	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-queries	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.2.5

Project	groupId	artifactId	version
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.2.5
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.2.5
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.2.5
	org.apache.solr	solr-cell	4.4.0-cdh5.2.5
	org.apache.solr	solr-clustering	4.4.0-cdh5.2.5
	org.apache.solr	solr-core	4.4.0-cdh5.2.5
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.2.5
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.2.5
	org.apache.solr	solr-langid	4.4.0-cdh5.2.5
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.2.5
	org.apache.solr	solr-solrj	4.4.0-cdh5.2.5
	org.apache.solr	solr-test-framework	4.4.0-cdh5.2.5
	org.apache.solr	solr-uima	4.4.0-cdh5.2.5
	org.apache.solr	solr-velocity	4.4.0-cdh5.2.5
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.2.5
	com.cloudera.search	search-mr	1.0.0-cdh5.2.5
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.2.5
	com.ngdata	hbase-indexer-cli	1.5-cdh5.2.5
	com.ngdata	hbase-indexer-common	1.5-cdh5.2.5
	com.ngdata	hbase-indexer-demo	1.5-cdh5.2.5
	com.ngdata	hbase-indexer-dist	1.5-cdh5.2.5
	com.ngdata	hbase-indexer-engine	1.5-cdh5.2.5
	com.ngdata	hbase-indexer-model	1.5-cdh5.2.5
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.2.5
	com.ngdata	hbase-indexer-mr	1.5-cdh5.2.5
	com.ngdata	hbase-indexer-server	1.5-cdh5.2.5
	com.ngdata	hbase-sep-api	1.5-cdh5.2.5
	com.ngdata	hbase-sep-demo	1.5-cdh5.2.5
	com.ngdata	hbase-sep-impl	1.5-cdh5.2.5

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.2.5
	com.ngdata	hbase-sep-tools	1.5-cdh5.2.5

CDH 5.2.4

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.2.4

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.2.4
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.2.4
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	250-mr1-cdh5.2.4
	org.apache.hadoop	hadoop-core	250-mr1-cdh5.2.4
	org.apache.hadoop	hadoop-examples	250-mr1-cdh5.2.4
	org.apache.hadoop	hadoop-minicluster	250-mr1-cdh5.2.4
	org.apache.hadoop	hadoop-streaming	250-mr1-cdh5.2.4
	org.apache.hadoop	hadoop-test	250-mr1-cdh5.2.4
	org.apache.hadoop	hadoop-tools	250-mr1-cdh5.2.4
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.2.4
	org.apache.hive	hive-beeline	0.13.1-cdh5.2.4
	org.apache.hive	hive-cli	0.13.1-cdh5.2.4
	org.apache.hive	hive-common	0.13.1-cdh5.2.4
	org.apache.hive	hive-contrib	0.13.1-cdh5.2.4
	org.apache.hive	hive-exec	0.13.1-cdh5.2.4
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.2.4
	org.apache.hive	hive-hwi	0.13.1-cdh5.2.4
	org.apache.hive	hive-jdbc	0.13.1-cdh5.2.4
	org.apache.hive	hive-metastore	0.13.1-cdh5.2.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-serde	0.13.1-cdh5.2.4
	org.apache.hive	hive-service	0.13.1-cdh5.2.4
	org.apache.hive	hive-shims	0.13.1-cdh5.2.4
	org.apache.hive	hive-testutils	0.13.1-cdh5.2.4
	org.apache.hcatalog	hive-hcatalog-core	0.13.1-cdh5.2.4
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.2.4
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.13.1-cdh5.2.4
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.2.4
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.2.4
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.2.4
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.2.4
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.2.4
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.2.4
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.2.4
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.2.4
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-common	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-examples	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-it	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-server	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-shell	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.2.4
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.2.4
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.2.4
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.2.4
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.2.4

Project	groupId	artifactId	version
	org.apache.pig	piggybank	0.12.0-cdh5.2.4
	org.apache.pig	pigsmoke	0.12.0-cdh5.2.4
	org.apache.pig	pigunit	0.12.0-cdh5.2.4
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-core	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-node	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.2.4
	org.apache.flume	flume-tools	1.5.0-cdh5.2.4
	org.apache.flume	flume-file-channel	1.5.0-cdh5.2.4
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.2.4
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.2.4
	org.apache.flume	flume-avro-source	1.5.0-cdh5.2.4
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.2.4
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.2.4
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.2.4
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.2.4
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.2.4
	org.apache.flume	flume-jms-source	1.5.0-cdh5.2.4
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.2.4
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.2.4
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.2.4
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-core	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-examples	4.0.0-cdh5.2.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-tools	4.0.0-cdh5.2.4
	org.apache.oozie	oozie-mini	4.0.0-cdh5.2.4
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.2.4
	org.apache.mahout	mahout-core	0.9-cdh5.2.4
	org.apache.mahout	mahout-examples	0.9-cdh5.2.4
	org.apache.mahout	mahout-integration	0.9-cdh5.2.4
	org.apache.mahout	mahout-math	0.9-cdh5.2.4
	org.apache.mahout	mahout-math-scala	0.9-cdh5.2.4
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-chef	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-cli	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-core	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-druid	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.2.4

Project	groupId	artifactId	version
	org.apache.whirr	whirr-examples	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-hama	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-pig	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-solr	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.2.4
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.2.4
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.2.4
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.2.4
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.2.4
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-dist	1.4.0-cdh5.2.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.2.4
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.2.4
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.2.4
	com.twitter	parquet-cascading	1.5.0-cdh5.2.4
	com.twitter	parquet-column	1.5.0-cdh5.2.4
	com.twitter	parquet-common	1.5.0-cdh5.2.4
	com.twitter	parquet-encoding	1.5.0-cdh5.2.4
	com.twitter	parquet-format	2.1.0-cdh5.2.4
	com.twitter	parquet-generator	1.5.0-cdh5.2.4
	com.twitter	parquet-hadoop	1.5.0-cdh5.2.4
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.2.4
	com.twitter	parquet-jackson	1.5.0-cdh5.2.4
	com.twitter	parquet-pig	1.5.0-cdh5.2.4
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.2.4
	com.twitter	parquet-protobuf	1.5.0-cdh5.2.4
	com.twitter	parquet-scrooge	1.5.0-cdh5.2.4
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.2.4
	com.twitter	parquet-thrift	1.5.0-cdh5.2.4
	com.twitter	parquet-tools	1.5.0-cdh5.2.4
Llama	com.doudera.llama	llama	1.0.0-cdh5.2.4
Apache Spark	org.apache.spark	spark-bagel_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-catalyst_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-core_2.10	1.1.0-cdh5.2.4

Project	groupId	artifactId	version
	org.apache.spark	spark-graphx_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-hive_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-mllib_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-sql_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-streaming-flume-sink_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-streaming-flume_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-streaming-kafka_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-streaming-mqtt_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-streaming-twitter_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-streaming-zeromq_2.10	1.1.0-cdh5.2.4
	org.apache.spark	spark-streaming_2.10	1.1.0-cdh5.2.4
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.2.4
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.2.4
	org.apache.crunch	crunch-core	0.11.0-cdh5.2.4
	org.apache.crunch	crunch-examples	0.11.0-cdh5.2.4
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.2.4
	org.apache.crunch	crunch-hive	0.11.0-cdh5.2.4
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.2.4
	org.apache.crunch	crunch-spark	0.11.0-cdh5.2.4
	org.apache.crunch	crunch-test	0.11.0-cdh5.2.4
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.2.4
	org.apache.avro	avro-compiler	1.7.6-cdh5.2.4
	org.apache.avro	avro-ipc	1.7.6-cdh5.2.4
	org.apache.avro	avro-mapred	1.7.6-cdh5.2.4
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.2.4
	org.apache.avro	avro-protobuf	1.7.6-cdh5.2.4
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.2.4
	org.apache.avro	avro-thrift	1.7.6-cdh5.2.4
	org.apache.avro	avro-tools	1.7.6-cdh5.2.4
	org.apache.avro	trevni-avro	1.7.6-cdh5.2.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	trevni-core	1.7.6-cdh5.2.4
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.2.4
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.2.4
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.2.4
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.2.4
	org.kitesdk	kite-data-hive	0.15.0-cdh5.2.4
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.2.4
	org.kitesdk	kite-data-spark	0.15.0-cdh5.2.4
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.2.4
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.2.4
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.2.4
	org.kitesdk	kite-tools	0.15.0-cdh5.2.4
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.2.4

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-classification	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-core	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-demo	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-facet	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-join	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-memory	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-misc	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-queries	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.2.4
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.2.4
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.2.4
	org.apache.solr	solr-cell	4.4.0-cdh5.2.4
	org.apache.solr	solr-clustering	4.4.0-cdh5.2.4
	org.apache.solr	solr-core	4.4.0-cdh5.2.4
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.2.4
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.2.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-langid	4.4.0-cdh5.2.4
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.2.4
	org.apache.solr	solr-solrj	4.4.0-cdh5.2.4
	org.apache.solr	solr-test-framework	4.4.0-cdh5.2.4
	org.apache.solr	solr-uima	4.4.0-cdh5.2.4
	org.apache.solr	solr-velocity	4.4.0-cdh5.2.4
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.2.4
	com.cloudera.search	search-mr	1.0.0-cdh5.2.4
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.2.4
	com.ngdata	hbase-indexer-cli	1.5-cdh5.2.4
	com.ngdata	hbase-indexer-common	1.5-cdh5.2.4
	com.ngdata	hbase-indexer-demo	1.5-cdh5.2.4
	com.ngdata	hbase-indexer-dist	1.5-cdh5.2.4
	com.ngdata	hbase-indexer-engine	1.5-cdh5.2.4
	com.ngdata	hbase-indexer-model	1.5-cdh5.2.4
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.2.4
	com.ngdata	hbase-indexer-mr	1.5-cdh5.2.4
	com.ngdata	hbase-indexer-server	1.5-cdh5.2.4
	com.ngdata	hbase-sep-api	1.5-cdh5.2.4
	com.ngdata	hbase-sep-demo	1.5-cdh5.2.4
	com.ngdata	hbase-sep-impl	1.5-cdh5.2.4
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.2.4
	com.ngdata	hbase-sep-tools	1.5-cdh5.2.4

CDH 5.2.3



Note: There is no CDH 5.0.7, 5.1.1, 5.1.6, 5.2.2, 5.2.7, 5.3.7, or 5.4.6 release.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.2.3

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.2.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.2.3
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-core	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-examples	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-minicluster	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-streaming	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-test	2.5.0-cdh5.2.3
	org.apache.hadoop	hadoop-tools	2.5.0-cdh5.2.3
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.2.3
	org.apache.hive	hive-beeline	0.13.1-cdh5.2.3
	org.apache.hive	hive-cli	0.13.1-cdh5.2.3
	org.apache.hive	hive-common	0.13.1-cdh5.2.3
	org.apache.hive	hive-contrib	0.13.1-cdh5.2.3
	org.apache.hive	hive-exec	0.13.1-cdh5.2.3
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.2.3
	org.apache.hive	hive-hwi	0.13.1-cdh5.2.3
	org.apache.hive	hive-jdbc	0.13.1-cdh5.2.3
	org.apache.hive	hive-metastore	0.13.1-cdh5.2.3
	org.apache.hive	hive-serde	0.13.1-cdh5.2.3
	org.apache.hive	hive-service	0.13.1-cdh5.2.3
	org.apache.hive	hive-shims	0.13.1-cdh5.2.3
	org.apache.hive	hive-testutils	0.13.1-cdh5.2.3
	org.apache.hive	hive-hcatalog-core	0.13.1-cdh5.2.3
	org.apache.hive	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.2.3
	org.apache.hive	hive-hcatalog-pig-adapter	0.13.1-cdh5.2.3

Project	groupId	artifactId	version
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.2.3
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.2.3
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.2.3
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.2.3
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.2.3
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.2.3
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.2.3
	org.apache.hiveshims	hive-shims-scheduler	0.13.1-cdh5.2.3
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-common	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-examples	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-it	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-server	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-shell	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.2.3
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.2.3
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.2.3
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.2.3
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.2.3
	org.apache.pig	piggybank	0.12.0-cdh5.2.3
	org.apache.pig	pigsmoke	0.12.0-cdh5.2.3
	org.apache.pig	pigunit	0.12.0-cdh5.2.3
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.2.3
	org.apache.flume	flume-ng-core	1.5.0-cdh5.2.3
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.2.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-node	1.5.0-cdh5.2.3
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.2.3
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.2.3
	org.apache.flume	flume-tools	1.5.0-cdh5.2.3
	org.apache.flume	flume-file-channel	1.5.0-cdh5.2.3
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.2.3
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.2.3
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.2.3
	org.apache.flume	flume-avro-source	1.5.0-cdh5.2.3
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.2.3
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.2.3
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.2.3
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.2.3
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.2.3
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.2.3
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.2.3
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.2.3
	org.apache.flume	flume-jms-source	1.5.0-cdh5.2.3
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.2.3
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.2.3
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.2.3
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-core	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-examples	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.2.3

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.2.3
	org.apache.oozie	oozie-tools	4.0.0-cdh5.2.3
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.2.3
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.2.3
	org.apache.mahout	mahout-core	0.9-cdh5.2.3
	org.apache.mahout	mahout-examples	0.9-cdh5.2.3
	org.apache.mahout	mahout-integration	0.9-cdh5.2.3
	org.apache.mahout	mahout-math	0.9-cdh5.2.3
	org.apache.mahout	mahout-math-scala	0.9-cdh5.2.3
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-chef	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-cli	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-core	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-druid	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-examples	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-hama	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.2.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-pig	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-solr	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.2.3
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.2.3
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.2.3
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.2.3
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.2.3
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-dist	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.2.3

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.2.3
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.2.3
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.2.3
	com.twitter	parquet-cascading	1.5.0-cdh5.2.3
	com.twitter	parquet-column	1.5.0-cdh5.2.3
	com.twitter	parquet-common	1.5.0-cdh5.2.3
	com.twitter	parquet-encoding	1.5.0-cdh5.2.3
	com.twitter	parquet-format	2.1.0-cdh5.2.3
	com.twitter	parquet-generator	1.5.0-cdh5.2.3
	com.twitter	parquet-hadoop	1.5.0-cdh5.2.3
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.2.3
	com.twitter	parquet-jackson	1.5.0-cdh5.2.3
	com.twitter	parquet-pig	1.5.0-cdh5.2.3
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.2.3
	com.twitter	parquet-protobuf	1.5.0-cdh5.2.3
	com.twitter	parquet-scrooge	1.5.0-cdh5.2.3
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.2.3
	com.twitter	parquet-thrift	1.5.0-cdh5.2.3
	com.twitter	parquet-tools	1.5.0-cdh5.2.3
Llama	com.doudera.llama	llama	1.0.0-cdh5.2.3
Apache Spark	org.apache.spark	spark-bagel_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-catalyst_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-core_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-graphx_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-hive_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-mllib_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-sql_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-streaming-flume-sink_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-streaming-flume_2.10	1.1.0-cdh5.2.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-kafka_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-streaming-mqtt_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-streaming-twitter_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-streaming-zeromq_2.10	1.1.0-cdh5.2.3
	org.apache.spark	spark-streaming_2.10	1.1.0-cdh5.2.3
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.2.3
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.2.3
	org.apache.crunch	crunch-core	0.11.0-cdh5.2.3
	org.apache.crunch	crunch-examples	0.11.0-cdh5.2.3
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.2.3
	org.apache.crunch	crunch-hive	0.11.0-cdh5.2.3
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.2.3
	org.apache.crunch	crunch-spark	0.11.0-cdh5.2.3
	org.apache.crunch	crunch-test	0.11.0-cdh5.2.3
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.2.3
	org.apache.avro	avro-compiler	1.7.6-cdh5.2.3
	org.apache.avro	avro-ipc	1.7.6-cdh5.2.3
	org.apache.avro	avro-mapred	1.7.6-cdh5.2.3
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.2.3
	org.apache.avro	avro-protobuf	1.7.6-cdh5.2.3
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.2.3
	org.apache.avro	avro-thrift	1.7.6-cdh5.2.3
	org.apache.avro	avro-tools	1.7.6-cdh5.2.3
	org.apache.avro	trevni-avro	1.7.6-cdh5.2.3
	org.apache.avro	trevni-core	1.7.6-cdh5.2.3
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.2.3
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.2.3
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.2.3
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.2.3
	org.kitesdk	kite-data-hive	0.15.0-cdh5.2.3

Project	groupId	artifactId	version
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.2.3
	org.kitesdk	kite-data-spark	0.15.0-cdh5.2.3
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.2.3
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.2.3
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.2.3
	org.kitesdk	kite-tools	0.15.0-cdh5.2.3
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.2.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-classification	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-core	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-demo	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-facet	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-join	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-memory	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-misc	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-queries	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.2.3
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.2.3
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.2.3
	org.apache.solr	solr-cell	4.4.0-cdh5.2.3
	org.apache.solr	solr-clustering	4.4.0-cdh5.2.3
	org.apache.solr	solr-core	4.4.0-cdh5.2.3
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.2.3
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.2.3
	org.apache.solr	solr-langid	4.4.0-cdh5.2.3
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.2.3
	org.apache.solr	solr-solrj	4.4.0-cdh5.2.3
	org.apache.solr	solr-test-framework	4.4.0-cdh5.2.3
	org.apache.solr	solr-uima	4.4.0-cdh5.2.3
	org.apache.solr	solr-velocity	4.4.0-cdh5.2.3

Project	groupId	artifactId	version
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.2.3
	com.cloudera.search	search-mr	1.0.0-cdh5.2.3
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.2.3
	com.ngdata	hbase-indexer-cli	1.5-cdh5.2.3
	com.ngdata	hbase-indexer-common	1.5-cdh5.2.3
	com.ngdata	hbase-indexer-demo	1.5-cdh5.2.3
	com.ngdata	hbase-indexer-dist	1.5-cdh5.2.3
	com.ngdata	hbase-indexer-engine	1.5-cdh5.2.3
	com.ngdata	hbase-indexer-model	1.5-cdh5.2.3
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.2.3
	com.ngdata	hbase-indexer-mr	1.5-cdh5.2.3
	com.ngdata	hbase-indexer-server	1.5-cdh5.2.3
	com.ngdata	hbase-sep-api	1.5-cdh5.2.3
	com.ngdata	hbase-sep-demo	1.5-cdh5.2.3
	com.ngdata	hbase-sep-impl	1.5-cdh5.2.3
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.2.3
	com.ngdata	hbase-sep-tools	1.5-cdh5.2.3

CDH 5.2.1

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.2.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.2.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-cdh5.2.1
	org.apache.hadoop	hadoop-core	2.5.0-cdh5.2.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-examples	2.5.0-mr1-cdh5.2.1
	org.apache.hadoop	hadoop-minicluster	2.5.0-mr1-cdh5.2.1
	org.apache.hadoop	hadoop-streaming	2.5.0-mr1-cdh5.2.1
	org.apache.hadoop	hadoop-test	2.5.0-mr1-cdh5.2.1
	org.apache.hadoop	hadoop-tools	2.5.0-mr1-cdh5.2.1
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.2.1
	org.apache.hive	hive-beeline	0.13.1-cdh5.2.1
	org.apache.hive	hive-cli	0.13.1-cdh5.2.1
	org.apache.hive	hive-common	0.13.1-cdh5.2.1
	org.apache.hive	hive-contrib	0.13.1-cdh5.2.1
	org.apache.hive	hive-exec	0.13.1-cdh5.2.1
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.2.1
	org.apache.hive	hive-hwi	0.13.1-cdh5.2.1
	org.apache.hive	hive-jdbc	0.13.1-cdh5.2.1
	org.apache.hive	hive-metastore	0.13.1-cdh5.2.1
	org.apache.hive	hive-serde	0.13.1-cdh5.2.1
	org.apache.hive	hive-service	0.13.1-cdh5.2.1
	org.apache.hive	hive-shims	0.13.1-cdh5.2.1
	org.apache.hive	hive-testutils	0.13.1-cdh5.2.1
	org.apache.hive	hive-hcatalog-core	0.13.1-cdh5.2.1
	org.apache.hive	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.2.1
	org.apache.hive	hive-hcatalog-pig-adapter	0.13.1-cdh5.2.1
	org.apache.hive	hive-hcatalog-server-extensions	0.13.1-cdh5.2.1
	org.apache.hive	hive-hcatalog-streaming	0.13.1-cdh5.2.1
	org.apache.hive	hive-webhcat	0.13.1-cdh5.2.1
	org.apache.hive	hive-webhcat-java-client	0.13.1-cdh5.2.1
	org.apache.hive	hive-shims-0.23	0.13.1-cdh5.2.1
	org.apache.hive	hive-shims-common	0.13.1-cdh5.2.1
	org.apache.hive	hive-shims-common-secure	0.13.1-cdh5.2.1
	org.apache.hive	hive-shims-scheduler	0.13.1-cdh5.2.1

Version and Download Information

Project	groupId	artifactId	version
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-common	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-examples	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-it	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-server	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-shell	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.2.1
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.2.1
	Apache ZooKeeper	org.apache.zookeeper	zookeeper
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.2.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.2.1
	org.apache.pig	piggybank	0.12.0-cdh5.2.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.2.1
	org.apache.pig	pigunit	0.12.0-cdh5.2.1
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-core	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-node	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.2.1
	org.apache.flume	flume-tools	1.5.0-cdh5.2.1
	org.apache.flume	flume-file-channel	1.5.0-cdh5.2.1
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.2.1
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.2.1

Project	groupId	artifactId	version
	org.apache.flume	flume-avro-source	1.5.0-cdh5.2.1
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.2.1
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.2.1
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.2.1
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.2.1
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.2.1
	org.apache.flume	flume-jms-source	1.5.0-cdh5.2.1
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.2.1
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.2.1
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.2.1
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-core	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-examples	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.2.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-tools	4.0.0-cdh5.2.1
	org.apache.oozie	oozie-test	4.0.0-cdh5.2.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.2.1
	org.apache.mahout	mahout-core	0.9-cdh5.2.1
	org.apache.mahout	mahout-examples	0.9-cdh5.2.1
	org.apache.mahout	mahout-integration	0.9-cdh5.2.1
	org.apache.mahout	mahout-math	0.9-cdh5.2.1
	org.apache.mahout	mahout-math-scala	0.9-cdh5.2.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.2.1
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.2.1
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.2.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.2.1

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.2.1
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.2.1
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.2.1
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.2.1
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.2.1
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.2.1
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.2.1
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.2.1
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.2.1
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.2.1
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-dist	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.2.1
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.2.1
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.2.1
	com.twitter	parquet-cascading	1.5.0-cdh5.2.1
	com.twitter	parquet-column	1.5.0-cdh5.2.1
	com.twitter	parquet-common	1.5.0-cdh5.2.1
	com.twitter	parquet-encoding	1.5.0-cdh5.2.1

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-format	2.1.0-cdh5.2.1
	com.twitter	parquet-generator	1.5.0-cdh5.2.1
	com.twitter	parquet-hadoop	1.5.0-cdh5.2.1
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.2.1
	com.twitter	parquet-jackson	1.5.0-cdh5.2.1
	com.twitter	parquet-pig	1.5.0-cdh5.2.1
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.2.1
	com.twitter	parquet-protobuf	1.5.0-cdh5.2.1
	com.twitter	parquet-scrooge	1.5.0-cdh5.2.1
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.2.1
	com.twitter	parquet-thrift	1.5.0-cdh5.2.1
	com.twitter	parquet-tools	1.5.0-cdh5.2.1
Llama	com.doudera.llama	llama	1.0.0-cdh5.2.1
Apache Spark	org.apache.spark	spark-bagel_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-catalyst_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-core_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-graphx_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-hive_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-mllib_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-sql_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-streaming-flume-sink_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-streaming-flume_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-streaming-kafka_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-streaming-mqtt_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-streaming-twitter_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-streaming-zeromq_2.10	1.1.0-cdh5.2.1
	org.apache.spark	spark-streaming_2.10	1.1.0-cdh5.2.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.2.1
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.2.1
	org.apache.crunch	crunch-core	0.11.0-cdh5.2.1

Project	groupId	artifactId	version
	org.apache.crunch	crunch-examples	0.11.0-cdh5.2.1
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.2.1
	org.apache.crunch	crunch-hive	0.11.0-cdh5.2.1
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.2.1
	org.apache.crunch	crunch-spark	0.11.0-cdh5.2.1
	org.apache.crunch	crunch-test	0.11.0-cdh5.2.1
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.2.1
	org.apache.avro	avro-compiler	1.7.6-cdh5.2.1
	org.apache.avro	avro-ipc	1.7.6-cdh5.2.1
	org.apache.avro	avro-mapred	1.7.6-cdh5.2.1
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.2.1
	org.apache.avro	avro-protobuf	1.7.6-cdh5.2.1
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.2.1
	org.apache.avro	avro-thrift	1.7.6-cdh5.2.1
	org.apache.avro	avro-tools	1.7.6-cdh5.2.1
	org.apache.avro	trevni-avro	1.7.6-cdh5.2.1
	org.apache.avro	trevni-core	1.7.6-cdh5.2.1
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.2.1
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.2.1
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.2.1
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.2.1
	org.kitesdk	kite-data-hive	0.15.0-cdh5.2.1
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.2.1
	org.kitesdk	kite-data-spark	0.15.0-cdh5.2.1
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.2.1
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.2.1

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.2.1
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.2.1
	org.kitesdk	kite-tools	0.15.0-cdh5.2.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-classification	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-core	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-demo	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-facet	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.2.1

Project	groupId	artifactId	version
	org.apache.lucene	lucene-join	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-memory	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-misc	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-queries	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.2.1
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.2.1
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.2.1
	org.apache.solr	solr-cell	4.4.0-cdh5.2.1
	org.apache.solr	solr-clustering	4.4.0-cdh5.2.1
	org.apache.solr	solr-core	4.4.0-cdh5.2.1
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.2.1
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.2.1
	org.apache.solr	solr-langid	4.4.0-cdh5.2.1
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.2.1
	org.apache.solr	solr-solrj	4.4.0-cdh5.2.1
	org.apache.solr	solr-test-framework	4.4.0-cdh5.2.1
	org.apache.solr	solr-uima	4.4.0-cdh5.2.1
	org.apache.solr	solr-velocity	4.4.0-cdh5.2.1
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.2.1
	com.cloudera.search	search-mr	1.0.0-cdh5.2.1
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.2.1
	com.ngdata	hbase-indexer-cli	1.5-cdh5.2.1
	com.ngdata	hbase-indexer-common	1.5-cdh5.2.1
	com.ngdata	hbase-indexer-demo	1.5-cdh5.2.1
	com.ngdata	hbase-indexer-dist	1.5-cdh5.2.1
	com.ngdata	hbase-indexer-engine	1.5-cdh5.2.1

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-model	1.5-cdh5.2.1
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.2.1
	com.ngdata	hbase-indexer-mr	1.5-cdh5.2.1
	com.ngdata	hbase-indexer-server	1.5-cdh5.2.1
	com.ngdata	hbase-sep-api	1.5-cdh5.2.1
	com.ngdata	hbase-sep-demo	1.5-cdh5.2.1
	com.ngdata	hbase-sep-impl	1.5-cdh5.2.1
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.2.1
	com.ngdata	hbase-sep-tools	1.5-cdh5.2.1

CDH 5.2.0

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-archives	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-assemblies	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-auth	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-common	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-datajoin	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-distcp	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-extras	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-gridmix	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-hdfs	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	2.5.0-cdh5.2.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-maven-plugins	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-minikdc	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-nfs	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-openstack	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-rumen	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-sls	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-api	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-client	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-common	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-server-common	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.5.0-cdh5.2.0
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.5.0-cdh5.2.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.5.0-mr1-cdh5.2.0
	org.apache.hadoop	hadoop-core	2.5.0-mr1-cdh5.2.0
	org.apache.hadoop	hadoop-examples	2.5.0-mr1-cdh5.2.0
	org.apache.hadoop	hadoop-minicluster	2.5.0-mr1-cdh5.2.0
	org.apache.hadoop	hadoop-streaming	2.5.0-mr1-cdh5.2.0
	org.apache.hadoop	hadoop-test	2.5.0-mr1-cdh5.2.0
	org.apache.hadoop	hadoop-tools	2.5.0-mr1-cdh5.2.0
Apache Hive	org.apache.hive	hive-ant	0.13.1-cdh5.2.0
	org.apache.hive	hive-beeline	0.13.1-cdh5.2.0
	org.apache.hive	hive-cli	0.13.1-cdh5.2.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-common	0.13.1-cdh5.2.0
	org.apache.hive	hive-contrib	0.13.1-cdh5.2.0
	org.apache.hive	hive-exec	0.13.1-cdh5.2.0
	org.apache.hive	hive-hbase-handler	0.13.1-cdh5.2.0
	org.apache.hive	hive-hwi	0.13.1-cdh5.2.0
	org.apache.hive	hive-jdbc	0.13.1-cdh5.2.0
	org.apache.hive	hive-metastore	0.13.1-cdh5.2.0
	org.apache.hive	hive-serde	0.13.1-cdh5.2.0
	org.apache.hive	hive-service	0.13.1-cdh5.2.0
	org.apache.hive	hive-shims	0.13.1-cdh5.2.0
	org.apache.hive	hive-testutils	0.13.1-cdh5.2.0
	org.apache.hcatalog	hive-hcatalog-core	0.13.1-cdh5.2.0
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.13.1-cdh5.2.0
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.13.1-cdh5.2.0
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.13.1-cdh5.2.0
	org.apache.hcatalog	hive-hcatalog-streaming	0.13.1-cdh5.2.0
	org.apache.hcatalog	hive-webhcat	0.13.1-cdh5.2.0
	org.apache.hcatalog	hive-webhcat-java-client	0.13.1-cdh5.2.0
	org.apache.hiveshims	hive-shims-0.23	0.13.1-cdh5.2.0
	org.apache.hiveshims	hive-shims-common	0.13.1-cdh5.2.0
	org.apache.hiveshims	hive-shims-common-secure	0.13.1-cdh5.2.0
Apache HBase	org.apache.hbase	hbase-client	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-common	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-examples	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-hadoop-compat	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-hadoop2-compat	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-it	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-prefix-tree	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-protocol	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-server	0.98.6-cdh5.2.0

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shell	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-testing-util	0.98.6-cdh5.2.0
	org.apache.hbase	hbase-thrift	0.98.6-cdh5.2.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.2.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.5-cdh5.2.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.2.0
	org.apache.pig	piggybank	0.12.0-cdh5.2.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.2.0
	org.apache.pig	pigunit	0.12.0-cdh5.2.0
Apache Flume	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-core	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-node	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.2.0
	org.apache.flume	flume-tools	1.5.0-cdh5.2.0
	org.apache.flume	flume-file-channel	1.5.0-cdh5.2.0
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.2.0
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.2.0
	org.apache.flume	flume-avro-source	1.5.0-cdh5.2.0
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.2.0
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.2.0
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.2.0
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-kafka-sink	1.5.0-cdh5.2.0
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.2.0
	org.apache.flume	flume-jms-source	1.5.0-cdh5.2.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-kafka-source	1.5.0-cdh5.2.0
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.2.0
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.2.0
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-core	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-examples	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-tools	4.0.0-cdh5.2.0
	org.apache.oozie	oozie-mini	4.0.0-cdh5.2.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.2.0
	org.apache.mahout	mahout-core	0.9-cdh5.2.0
	org.apache.mahout	mahout-examples	0.9-cdh5.2.0
	org.apache.mahout	mahout-integration	0.9-cdh5.2.0
	org.apache.mahout	mahout-math	0.9-cdh5.2.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.2.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.2.0

Project	groupId	artifactId	version
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.2.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.2.0
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.2.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.2.0
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.2.0

Version and Download Information

Project	groupId	artifactId	version
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-binding-solr	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-core-common	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-core-model-db	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-core-model-search	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-dist	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-policy-common	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-policy-db	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-policy-search	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-provider-cache	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-provider-common	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-provider-db	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-provider-file	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-tests-hive	1.4.0-cdh5.2.0
	org.apache.sentry	sentry-tests-solr	1.4.0-cdh5.2.0
Parquet	com.twitter	parquet-avro	1.5.0-cdh5.2.0
	com.twitter	parquet-cascading	1.5.0-cdh5.2.0
	com.twitter	parquet-column	1.5.0-cdh5.2.0
	com.twitter	parquet-common	1.5.0-cdh5.2.0
	com.twitter	parquet-encoding	1.5.0-cdh5.2.0
	com.twitter	parquet-format	2.1.0-cdh5.2.0
	com.twitter	parquet-generator	1.5.0-cdh5.2.0
	com.twitter	parquet-hadoop	1.5.0-cdh5.2.0
	com.twitter	parquet-hadoop-bundle	1.5.0-cdh5.2.0
	com.twitter	parquet-jackson	1.5.0-cdh5.2.0
	com.twitter	parquet-pig	1.5.0-cdh5.2.0
	com.twitter	parquet-pig-bundle	1.5.0-cdh5.2.0
	com.twitter	parquet-protobuf	1.5.0-cdh5.2.0
	com.twitter	parquet-scrooge	1.5.0-cdh5.2.0
	com.twitter	parquet-test-hadoop2	1.5.0-cdh5.2.0

Project	groupId	artifactId	version
	com.twitter	parquet-thrift	1.5.0-cdh5.2.0
	com.twitter	parquet-tools	1.5.0-cdh5.2.0
Llama	com.doudera.llama	llama	1.0.0-cdh5.2.0
Apache Spark	org.apache.spark	spark-bagel_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-catalyst_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-core_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-graphx_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-hive_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-mllib_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-sql_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-streaming-flume-sink_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-streaming-flume_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-streaming-kafka_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-streaming-twitter_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.1.0-cdh5.2.0
	org.apache.spark	spark-streaming_2.10	1.1.0-cdh5.2.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0-cdh5.2.0
	org.apache.crunch	crunch-contrib	0.11.0-cdh5.2.0
	org.apache.crunch	crunch-core	0.11.0-cdh5.2.0
	org.apache.crunch	crunch-examples	0.11.0-cdh5.2.0
	org.apache.crunch	crunch-hbase	0.11.0-cdh5.2.0
	org.apache.crunch	crunch-hive	0.11.0-cdh5.2.0
	org.apache.crunch	crunch-scrunch	0.11.0-cdh5.2.0
	org.apache.crunch	crunch-spark	0.11.0-cdh5.2.0
	org.apache.crunch	crunch-test	0.11.0-cdh5.2.0
Apache Avro	org.apache.avro	avro	1.7.6-cdh5.2.0
	org.apache.avro	avro-compiler	1.7.6-cdh5.2.0
	org.apache.avro	avro-ipc	1.7.6-cdh5.2.0
	org.apache.avro	avro-mapred	1.7.6-cdh5.2.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.avro	avro-maven-plugin	1.7.6-cdh5.2.0
	org.apache.avro	avro-protobuf	1.7.6-cdh5.2.0
	org.apache.avro	avro-service-archetype	1.7.6-cdh5.2.0
	org.apache.avro	avro-thrift	1.7.6-cdh5.2.0
	org.apache.avro	avro-tools	1.7.6-cdh5.2.0
	org.apache.avro	trevni-avro	1.7.6-cdh5.2.0
	org.apache.avro	trevni-core	1.7.6-cdh5.2.0
Kite SDK	org.kitesdk	kite-data-core	0.15.0-cdh5.2.0
	org.kitesdk	kite-data-crunch	0.15.0-cdh5.2.0
	org.kitesdk	kite-data-hbase	0.15.0-cdh5.2.0
	org.kitesdk	kite-data-hcatalog	0.15.0-cdh5.2.0
	org.kitesdk	kite-data-hive	0.15.0-cdh5.2.0
	org.kitesdk	kite-data-mapreduce	0.15.0-cdh5.2.0
	org.kitesdk	kite-data-spark	0.15.0-cdh5.2.0
	org.kitesdk	kite-hadoop-compatibility	0.15.0-cdh5.2.0
	org.kitesdk	kite-maven-plugin	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-avro	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-core	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-hadoop-core	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-json	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-maxmind	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-metrics-servlets	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-protobuf	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-saxon	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-solr-cell	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-solr-core	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-tika-core	0.15.0-cdh5.2.0

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-tika-decompress	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-twitter	0.15.0-cdh5.2.0
	org.kitesdk	kite-morphlines-useragent	0.15.0-cdh5.2.0
	org.kitesdk	kite-tools	0.15.0-cdh5.2.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-classification	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-core	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-demo	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-facet	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-join	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-memory	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-misc	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-queries	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.2.0
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.2.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.2.0
	org.apache.solr	solr-cell	4.4.0-cdh5.2.0
	org.apache.solr	solr-clustering	4.4.0-cdh5.2.0
	org.apache.solr	solr-core	4.4.0-cdh5.2.0
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.2.0
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.2.0
	org.apache.solr	solr-langid	4.4.0-cdh5.2.0
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.2.0
	org.apache.solr	solr-solrj	4.4.0-cdh5.2.0
	org.apache.solr	solr-test-framework	4.4.0-cdh5.2.0
	org.apache.solr	solr-uima	4.4.0-cdh5.2.0
	org.apache.solr	solr-velocity	4.4.0-cdh5.2.0
Cloudera Search	com.cloudera.search	search-crunch	1.0.0-cdh5.2.0
	com.cloudera.search	search-mr	1.0.0-cdh5.2.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.2.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.2.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.2.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.2.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.2.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.2.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.2.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.2.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.2.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.2.0
	com.ngdata	hbase-sep-api	1.5-cdh5.2.0
	com.ngdata	hbase-sep-demo	1.5-cdh5.2.0
	com.ngdata	hbase-sep-impl	1.5-cdh5.2.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.2.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.2.0

Maven Artifacts for CDH 5.1.x Releases
CDH 5.1.5

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.1.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.1.5
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-core	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-examples	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-minicluster	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-streaming	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-test	2.3.0-cdh5.1.5
	org.apache.hadoop	hadoop-tools	2.3.0-cdh5.1.5
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.1.5
	org.apache.hive	hive-beeline	0.12.0-cdh5.1.5
	org.apache.hive	hive-cli	0.12.0-cdh5.1.5
	org.apache.hive	hive-common	0.12.0-cdh5.1.5
	org.apache.hive	hive-contrib	0.12.0-cdh5.1.5
	org.apache.hive	hive-exec	0.12.0-cdh5.1.5
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.1.5
	org.apache.hive	hive-hwi	0.12.0-cdh5.1.5
	org.apache.hive	hive-jdbc	0.12.0-cdh5.1.5
	org.apache.hive	hive-metastore	0.12.0-cdh5.1.5
	org.apache.hive	hive-serde	0.12.0-cdh5.1.5
	org.apache.hive	hive-service	0.12.0-cdh5.1.5
	org.apache.hive	hive-shims	0.12.0-cdh5.1.5
	org.apache.hive	hive-testutils	0.12.0-cdh5.1.5
	org.apache.hive	hive-hcatalog-core	0.12.0-cdh5.1.5

Project	groupId	artifactId	version
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.1.5
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.1.5
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.1.5
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.1.5
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.1.5
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.1.5
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.1.5
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.1.5
Apache HBase	org.apache.hbase	hbase-client	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-common	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-examples	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-hadoop-compat	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-hadoop2-compat	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-it	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-prefix-tree	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-protocol	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-server	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-shell	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-testing-util	0.98.1-cdh5.1.5
	org.apache.hbase	hbase-thrift	0.98.1-cdh5.1.5
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.1.5
Apache Sqoop	org.apache.sqoop	sqoop	1.4.4-cdh5.1.5
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.1.5
	org.apache.pig	pigsmoke	0.12.0-cdh5.1.5
	org.apache.pig	pigunit	0.12.0-cdh5.1.5
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.1.5
	org.apache.flume	flume-ng-core	1.5.0-cdh5.1.5
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.1.5
	org.apache.flume	flume-ng-node	1.5.0-cdh5.1.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.1.5
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.1.5
	org.apache.flume	flume-tools	1.5.0-cdh5.1.5
	org.apache.flume	flume-file-channel	1.5.0-cdh5.1.5
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.1.5
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.1.5
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.1.5
	org.apache.flume	flume-avro-source	1.5.0-cdh5.1.5
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.1.5
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.1.5
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.1.5
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.1.5
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.1.5
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.1.5
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.1.5
	org.apache.flume	flume-jms-source	1.5.0-cdh5.1.5
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.1.5
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.1.5
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-core	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-examples	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-hadoop	20141115-cdh5.1.5
	org.apache.oozie	oozie-hadoop-distcp	20141115-cdh5.1.5
	org.apache.oozie	oozie-hadoop-test	20141115-cdh5.1.5
	org.apache.oozie	oozie-hadoop-utils	20141115-cdh5.1.5
	org.apache.oozie	oozie-hbase	081415-cdh5.1.5
	org.apache.oozie	oozie-hcatalog	021415-cdh5.1.5
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.1.5

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-tools	4.0.0-cdh5.1.5
	org.apache.oozie	oozie-test	4.0.0-cdh5.1.5
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.1.5
	org.apache.mahout	mahout-core	0.9-cdh5.1.5
	org.apache.mahout	mahout-examples	0.9-cdh5.1.5
	org.apache.mahout	mahout-integration	0.9-cdh5.1.5
	org.apache.mahout	mahout-math	0.9-cdh5.1.5
	org.apache.mahout	mahout-math-scala	0.9-cdh5.1.5
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-chef	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-cli	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-core	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-druid	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-examples	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-hama	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-pig	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.1.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-solr	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.1.5
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.1.5
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.1.5
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.1.5
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.1.5
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-binding-solr	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-core-common	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-core-model-db	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-core-model-search	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-dist	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-policy-common	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-policy-db	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-policy-search	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-provider-cache	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-provider-common	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-provider-db	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-provider-file	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-tests-hive	1.3.0-cdh5.1.5
	org.apache.sentry	sentry-tests-solr	1.3.0-cdh5.1.5

Project	groupId	artifactId	version
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.1.5
	com.twitter	parquet-cascading	1.2.5-cdh5.1.5
	com.twitter	parquet-column	1.2.5-cdh5.1.5
	com.twitter	parquet-common	1.2.5-cdh5.1.5
	com.twitter	parquet-encoding	1.2.5-cdh5.1.5
	com.twitter	parquet-format	1.0.0-cdh5.1.5
	com.twitter	parquet-generator	1.2.5-cdh5.1.5
	com.twitter	parquet-hadoop	1.2.5-cdh5.1.5
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.1.5
	com.twitter	parquet-pig	1.2.5-cdh5.1.5
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.1.5
	com.twitter	parquet-scrooge	1.2.5-cdh5.1.5
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.1.5
	com.twitter	parquet-thrift	1.2.5-cdh5.1.5
Llama	com.douderallama	llama	1.0.0-cdh5.1.5
Apache Spark	org.apache.spark	spark-bagel_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-catalyst_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-core_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-examples_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-graphx_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-hive_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-mllib_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-repl_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-sql_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-streaming-flume_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-streaming-kafka_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-streaming-mqtt_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-streaming-twitter_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-streaming-zeromq_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-streaming_2.10	1.0.0-cdh5.1.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-tools_2.10	1.0.0-cdh5.1.5
	org.apache.spark	spark-yarn_2.10	1.0.0-cdh5.1.5
Apache Crunch	org.apache.crunch	crunch-archetype	0.10.0-cdh5.1.5
	org.apache.crunch	crunch-contrib	0.10.0-cdh5.1.5
	org.apache.crunch	crunch-core	0.10.0-cdh5.1.5
	org.apache.crunch	crunch-examples	0.10.0-cdh5.1.5
	org.apache.crunch	crunch-hbase	0.10.0-cdh5.1.5
	org.apache.crunch	crunch-scrunch	0.10.0-cdh5.1.5
	org.apache.crunch	crunch-spark	0.10.0-cdh5.1.5
	org.apache.crunch	crunch-test	0.10.0-cdh5.1.5
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.1.5
	org.apache.avro	avro-compiler	1.7.5-cdh5.1.5
	org.apache.avro	avro-ipc	1.7.5-cdh5.1.5
	org.apache.avro	avro-mapred	1.7.5-cdh5.1.5
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.1.5
	org.apache.avro	avro-protobuf	1.7.5-cdh5.1.5
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.1.5
	org.apache.avro	avro-thrift	1.7.5-cdh5.1.5
	org.apache.avro	avro-tools	1.7.5-cdh5.1.5
	org.apache.avro	trevni-avro	1.7.5-cdh5.1.5
	org.apache.avro	trevni-core	1.7.5-cdh5.1.5
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.1.5
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.1.5
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.1.5
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.1.5
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.1.5

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.1.5
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.1.5
	org.kitesdk	kite-tools	0.10.0-cdh5.1.5
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-classification	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-core	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-demo	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-facet	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-join	4.4.0-cdh5.1.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-memory	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-misc	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-queries	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.1.5
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.1.5
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.1.5
	org.apache.solr	solr-cell	4.4.0-cdh5.1.5
	org.apache.solr	solr-clustering	4.4.0-cdh5.1.5
	org.apache.solr	solr-core	4.4.0-cdh5.1.5
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.1.5
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.1.5
	org.apache.solr	solr-langid	4.4.0-cdh5.1.5
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.1.5
	org.apache.solr	solr-solrj	4.4.0-cdh5.1.5
	org.apache.solr	solr-test-framework	4.4.0-cdh5.1.5
	org.apache.solr	solr-uima	4.4.0-cdh5.1.5
	org.apache.solr	solr-velocity	4.4.0-cdh5.1.5
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.1.5
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.1.5
	com.ngdata	hbase-indexer-cli	1.5-cdh5.1.5
	com.ngdata	hbase-indexer-common	1.5-cdh5.1.5
	com.ngdata	hbase-indexer-demo	1.5-cdh5.1.5
	com.ngdata	hbase-indexer-dist	1.5-cdh5.1.5
	com.ngdata	hbase-indexer-engine	1.5-cdh5.1.5
	com.ngdata	hbase-indexer-model	1.5-cdh5.1.5
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.1.5

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-mr	1.5-cdh5.1.5
	com.ngdata	hbase-indexer-server	1.5-cdh5.1.5
	com.ngdata	hbase-sep-api	1.5-cdh5.1.5
	com.ngdata	hbase-sep-demo	1.5-cdh5.1.5
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.1.5
	com.ngdata	hbase-sep-tools	1.5-cdh5.1.5

CDH 5.1.4

The following table lists the project name, groupId, artifactId, and version required to access each CDH artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.1.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.1.4
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.1.4
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.3.0-mr1-cdh5.1.4
	org.apache.hadoop	hadoop-core	2.3.0-mr1-cdh5.1.4
	org.apache.hadoop	hadoop-examples	2.3.0-mr1-cdh5.1.4
	org.apache.hadoop	hadoop-minicluster	2.3.0-mr1-cdh5.1.4
	org.apache.hadoop	hadoop-streaming	2.3.0-mr1-cdh5.1.4
	org.apache.hadoop	hadoop-test	2.3.0-mr1-cdh5.1.4
	org.apache.hadoop	hadoop-tools	2.3.0-mr1-cdh5.1.4
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.1.4
	org.apache.hive	hive-beeline	0.12.0-cdh5.1.4
	org.apache.hive	hive-cli	0.12.0-cdh5.1.4
	org.apache.hive	hive-common	0.12.0-cdh5.1.4
	org.apache.hive	hive-contrib	0.12.0-cdh5.1.4
	org.apache.hive	hive-exec	0.12.0-cdh5.1.4
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.1.4
	org.apache.hive	hive-hwi	0.12.0-cdh5.1.4

Project	groupId	artifactId	version
	org.apache.hive	hive-jdbc	0.12.0-cdh5.1.4
	org.apache.hive	hive-metastore	0.12.0-cdh5.1.4
	org.apache.hive	hive-serde	0.12.0-cdh5.1.4
	org.apache.hive	hive-service	0.12.0-cdh5.1.4
	org.apache.hive	hive-shims	0.12.0-cdh5.1.4
	org.apache.hive	hive-testutils	0.12.0-cdh5.1.4
	org.apache.hcatalog	hive-hcatalog-core	0.12.0-cdh5.1.4
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.1.4
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.1.4
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.1.4
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.1.4
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.1.4
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.1.4
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.1.4
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.1.4
Apache HBase	org.apache.hbase	hbase-client	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-common	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-examples	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-hadoop-compat	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-hadoop2-compat	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-it	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-prefix-tree	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-protocol	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-server	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-shell	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-testing-util	0.98.1-cdh5.1.4
	org.apache.hbase	hbase-thrift	0.98.1-cdh5.1.4
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.1.4
Apache Sqoop	org.apache.sqoop	connector-sdk	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop	1.4.4-cdh5.1.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.1.4
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.1.4
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.1.4
	org.apache.pig	pigsmoke	0.12.0-cdh5.1.4
	org.apache.pig	pigunit	0.12.0-cdh5.1.4
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.1.4
	org.apache.flume	flume-ng-core	1.5.0-cdh5.1.4
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.1.4
	org.apache.flume	flume-ng-node	1.5.0-cdh5.1.4
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.1.4
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.1.4
	org.apache.flume	flume-tools	1.5.0-cdh5.1.4
	org.apache.flume	flume-file-channel	1.5.0-cdh5.1.4
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.1.4
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.1.4
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.1.4
	org.apache.flume	flume-avro-source	1.5.0-cdh5.1.4
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.1.4
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.1.4
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.1.4
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.1.4
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.1.4

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.1.4
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.1.4
	org.apache.flume	flume-jms-source	1.5.0-cdh5.1.4
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.1.4
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.1.4
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-core	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-examples	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-tools	4.0.0-cdh5.1.4
	org.apache.oozie	oozie-mini	4.0.0-cdh5.1.4
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.1.4
	org.apache.mahout	mahout-core	0.9-cdh5.1.4
	org.apache.mahout	mahout-examples	0.9-cdh5.1.4
	org.apache.mahout	mahout-integration	0.9-cdh5.1.4
	org.apache.mahout	mahout-math	0.9-cdh5.1.4
	org.apache.mahout	mahout-math-scala	0.9-cdh5.1.4

Version and Download Information

Project	groupId	artifactId	version
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-chef	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-cli	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-core	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-druid	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-examples	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-hama	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-pig	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-solr	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.1.4
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.1.4
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.1.4
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-binding-solr	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-core-common	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-core-model-db	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-core-model-search	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-dist	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-policy-common	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-policy-db	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-policy-search	1.3.0-cdh5.1.4

Project	groupId	artifactId	version
	org.apache.sentry	sentry-provider-cache	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-provider-common	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-provider-db	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-provider-file	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-tests-hive	1.3.0-cdh5.1.4
	org.apache.sentry	sentry-tests-solr	1.3.0-cdh5.1.4
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.1.4
	com.twitter	parquet-cascading	1.2.5-cdh5.1.4
	com.twitter	parquet-column	1.2.5-cdh5.1.4
	com.twitter	parquet-common	1.2.5-cdh5.1.4
	com.twitter	parquet-encoding	1.2.5-cdh5.1.4
	com.twitter	parquet-format	1.0.0-cdh5.1.4
	com.twitter	parquet-generator	1.2.5-cdh5.1.4
	com.twitter	parquet-hadoop	1.2.5-cdh5.1.4
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.1.4
	com.twitter	parquet-pig	1.2.5-cdh5.1.4
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.1.4
	com.twitter	parquet-scrooge	1.2.5-cdh5.1.4
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.1.4
	com.twitter	parquet-thrift	1.2.5-cdh5.1.4
Llama	com.doudera.llama	llama	1.0.0-cdh5.1.4
Apache Spark	org.apache.spark	spark-bagel_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-catalyst_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-core_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-examples_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-graphx_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-hive_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-mllib_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-repl_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-sql_2.10	1.0.0-cdh5.1.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-streaming-flume_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-streaming-kafka_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-streaming-mqtt_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-streaming-twitter_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-streaming-zeromq_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-streaming_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-tools_2.10	1.0.0-cdh5.1.4
	org.apache.spark	spark-yarn_2.10	1.0.0-cdh5.1.4
Apache Crunch	org.apache.crunch	crunch-archetype	0.10.0-cdh5.1.4
	org.apache.crunch	crunch-contrib	0.10.0-cdh5.1.4
	org.apache.crunch	crunch-core	0.10.0-cdh5.1.4
	org.apache.crunch	crunch-examples	0.10.0-cdh5.1.4
	org.apache.crunch	crunch-hbase	0.10.0-cdh5.1.4
	org.apache.crunch	crunch-scrunch	0.10.0-cdh5.1.4
	org.apache.crunch	crunch-spark	0.10.0-cdh5.1.4
	org.apache.crunch	crunch-test	0.10.0-cdh5.1.4
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.1.4
	org.apache.avro	avro-compiler	1.7.5-cdh5.1.4
	org.apache.avro	avro-ipc	1.7.5-cdh5.1.4
	org.apache.avro	avro-mapred	1.7.5-cdh5.1.4
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.1.4
	org.apache.avro	avro-protobuf	1.7.5-cdh5.1.4
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.1.4
	org.apache.avro	avro-thrift	1.7.5-cdh5.1.4
	org.apache.avro	avro-tools	1.7.5-cdh5.1.4
	org.apache.avro	trevni-avro	1.7.5-cdh5.1.4
	org.apache.avro	trevni-core	1.7.5-cdh5.1.4
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.1.4
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.1.4
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.1.4

Project	groupId	artifactId	version
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.1.4
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.1.4
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.1.4
	org.kitesdk	kite-tools	0.10.0-cdh5.1.4
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-classification	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.1.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-core	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-demo	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-facet	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-join	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-memory	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-misc	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-queries	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.1.4
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.1.4
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.1.4
	org.apache.solr	solr-cell	4.4.0-cdh5.1.4
	org.apache.solr	solr-clustering	4.4.0-cdh5.1.4
	org.apache.solr	solr-core	4.4.0-cdh5.1.4
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.1.4
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.1.4
	org.apache.solr	solr-langid	4.4.0-cdh5.1.4
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.1.4
	org.apache.solr	solr-solrj	4.4.0-cdh5.1.4
	org.apache.solr	solr-test-framework	4.4.0-cdh5.1.4
	org.apache.solr	solr-uima	4.4.0-cdh5.1.4
	org.apache.solr	solr-velocity	4.4.0-cdh5.1.4
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.1.4
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.1.4
	com.ngdata	hbase-indexer-cli	1.5-cdh5.1.4

Project	groupId	artifactId	version
	com.ngdata	hbase-indexer-common	1.5-cdh5.1.4
	com.ngdata	hbase-indexer-demo	1.5-cdh5.1.4
	com.ngdata	hbase-indexer-dist	1.5-cdh5.1.4
	com.ngdata	hbase-indexer-engine	1.5-cdh5.1.4
	com.ngdata	hbase-indexer-model	1.5-cdh5.1.4
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.1.4
	com.ngdata	hbase-indexer-mr	1.5-cdh5.1.4
	com.ngdata	hbase-indexer-server	1.5-cdh5.1.4
	com.ngdata	hbase-sep-api	1.5-cdh5.1.4
	com.ngdata	hbase-sep-demo	1.5-cdh5.1.4
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.1.4
	com.ngdata	hbase-sep-tools	1.5-cdh5.1.4

CDH 5.1.3

The following table lists the project name, groupId, artifactId, and version required to access each CDH 5 artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.1.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.1.3
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.1.3
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.3.0-mr1-cdh5.1.3
	org.apache.hadoop	hadoop-core	2.3.0-mr1-cdh5.1.3
	org.apache.hadoop	hadoop-examples	2.3.0-mr1-cdh5.1.3
	org.apache.hadoop	hadoop-minicluster	2.3.0-mr1-cdh5.1.3
	org.apache.hadoop	hadoop-streaming	2.3.0-mr1-cdh5.1.3
	org.apache.hadoop	hadoop-test	2.3.0-mr1-cdh5.1.3
	org.apache.hadoop	hadoop-tools	2.3.0-mr1-cdh5.1.3
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.1.3
	org.apache.hive	hive-beeline	0.12.0-cdh5.1.3

Project	groupId	artifactId	version
	org.apache.hive	hive-cli	0.12.0-cdh5.1.3
	org.apache.hive	hive-common	0.12.0-cdh5.1.3
	org.apache.hive	hive-contrib	0.12.0-cdh5.1.3
	org.apache.hive	hive-exec	0.12.0-cdh5.1.3
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.1.3
	org.apache.hive	hive-hwi	0.12.0-cdh5.1.3
	org.apache.hive	hive-jdbc	0.12.0-cdh5.1.3
	org.apache.hive	hive-metastore	0.12.0-cdh5.1.3
	org.apache.hive	hive-serde	0.12.0-cdh5.1.3
	org.apache.hive	hive-service	0.12.0-cdh5.1.3
	org.apache.hive	hive-shims	0.12.0-cdh5.1.3
	org.apache.hive	hive-testutils	0.12.0-cdh5.1.3
	org.apache.hcatalog	hive-hcatalog-core	0.12.0-cdh5.1.3
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.1.3
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.1.3
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.1.3
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.1.3
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.1.3
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.1.3
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.1.3
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.1.3
Apache HBase	org.apache.hbase	hbase-client	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-common	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-examples	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-hadoop-compat	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-hadoop2-compat	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-it	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-prefix-tree	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-protocol	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-server	0.98.1-cdh5.1.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shell	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-testing-util	0.98.1-cdh5.1.3
	org.apache.hbase	hbase-thrift	0.98.1-cdh5.1.3
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.1.3
Apache Sqoop	org.apache.sqoop	sqoop	1.4.4-cdh5.1.3
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.1.3
	org.apache.pig	pigsmoke	0.12.0-cdh5.1.3
	org.apache.pig	pigunit	0.12.0-cdh5.1.3
Apache Flume	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.1.3
	org.apache.flume	flume-ng-core	1.5.0-cdh5.1.3
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.1.3
	org.apache.flume	flume-ng-node	1.5.0-cdh5.1.3
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.1.3
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.1.3
	org.apache.flume	flume-tools	1.5.0-cdh5.1.3
	org.apache.flume	flume-file-channel	1.5.0-cdh5.1.3
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.1.3
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.1.3
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.1.3
	org.apache.flume	flume-avro-source	1.5.0-cdh5.1.3
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.1.3
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.1.3
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.1.3
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.1.3
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.1.3
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.1.3
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.1.3
	org.apache.flume	flume-jms-source	1.5.0-cdh5.1.3
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.1.3
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.1.3

Project	groupId	artifactId	version
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-core	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-examples	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-tools	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-test	4.0.0-cdh5.1.3
	org.apache.oozie	oozie-mini	4.0.0-cdh5.1.3
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.1.3
	org.apache.mahout	mahout-core	0.9-cdh5.1.3
	org.apache.mahout	mahout-examples	0.9-cdh5.1.3
	org.apache.mahout	mahout-integration	0.9-cdh5.1.3
	org.apache.mahout	mahout-math	0.9-cdh5.1.3
	org.apache.mahout	mahout-math-scala	0.9-cdh5.1.3
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-chef	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-cli	0.9.0-cdh5.1.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.whirr	whirr-core	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-druid	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-examples	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-hama	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-pig	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-solr	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.1.3
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.1.3
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.1.3
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.1.3
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.1.3
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-binding-solr	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-core-common	1.3.0-cdh5.1.3

Project	groupId	artifactId	version
	org.apache.sentry	sentry-core-model-db	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-core-model-search	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-dist	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-policy-common	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-policy-db	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-policy-search	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-provider-cache	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-provider-common	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-provider-db	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-provider-file	1.3.0-cdh5.1.3
	org.apache.sentry	sentry-tests-hive	1.3.0-cdh5.1.3
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.1.3
	com.twitter	parquet-cascading	1.2.5-cdh5.1.3
	com.twitter	parquet-column	1.2.5-cdh5.1.3
	com.twitter	parquet-common	1.2.5-cdh5.1.3
	com.twitter	parquet-encoding	1.2.5-cdh5.1.3
	com.twitter	parquet-format	1.0.0-cdh5.1.3
	com.twitter	parquet-generator	1.2.5-cdh5.1.3
	com.twitter	parquet-hadoop	1.2.5-cdh5.1.3
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.1.3
	com.twitter	parquet-pig	1.2.5-cdh5.1.3
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.1.3
	com.twitter	parquet-scrooge	1.2.5-cdh5.1.3
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.1.3
	com.twitter	parquet-thrift	1.2.5-cdh5.1.3
Llama	com.doudera.llama	llama	1.0.0-cdh5.1.3
Apache Spark	org.apache.spark	spark-bagel_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-catalyst_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-core_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-examples_2.10	1.0.0-cdh5.1.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.spark	spark-graphx_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-hive_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-mllib_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-repl_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-sql_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-streaming-flume_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-streaming-kafka_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-streaming-mqtt_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-streaming-twitter_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-streaming-zeromq_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-streaming_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-tools_2.10	1.0.0-cdh5.1.3
	org.apache.spark	spark-yarn_2.10	1.0.0-cdh5.1.3
Apache Crunch	org.apache.crunch	crunch-archetype	0.10.0-cdh5.1.3
	org.apache.crunch	crunch-contrib	0.10.0-cdh5.1.3
	org.apache.crunch	crunch-core	0.10.0-cdh5.1.3
	org.apache.crunch	crunch-examples	0.10.0-cdh5.1.3
	org.apache.crunch	crunch-hbase	0.10.0-cdh5.1.3
	org.apache.crunch	crunch-scrunch	0.10.0-cdh5.1.3
	org.apache.crunch	crunch-spark	0.10.0-cdh5.1.3
	org.apache.crunch	crunch-test	0.10.0-cdh5.1.3
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.1.3
	org.apache.avro	avro-compiler	1.7.5-cdh5.1.3
	org.apache.avro	avro-ipc	1.7.5-cdh5.1.3
	org.apache.avro	avro-mapred	1.7.5-cdh5.1.3
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.1.3
	org.apache.avro	avro-protobuf	1.7.5-cdh5.1.3
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.1.3
	org.apache.avro	avro-thrift	1.7.5-cdh5.1.3
	org.apache.avro	avro-tools	1.7.5-cdh5.1.3

Project	groupId	artifactId	version
	org.apache.avro	trevni-avro	1.7.5-cdh5.1.3
	org.apache.avro	trevni-core	1.7.5-cdh5.1.3
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.1.3
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.1.3
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.1.3
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.1.3
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.1.3
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.1.3
	org.kitesdk	kite-tools	0.10.0-cdh5.1.3
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.1.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-classification	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-core	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-demo	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-facet	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-join	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-memory	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-misc	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-queries	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.1.3
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.1.3
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.1.3
	org.apache.solr	solr-cell	4.4.0-cdh5.1.3
	org.apache.solr	solr-clustering	4.4.0-cdh5.1.3
	org.apache.solr	solr-core	4.4.0-cdh5.1.3
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.1.3
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.1.3
	org.apache.solr	solr-langid	4.4.0-cdh5.1.3
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.1.3
	org.apache.solr	solr-solrj	4.4.0-cdh5.1.3
	org.apache.solr	solr-test-framework	4.4.0-cdh5.1.3

Project	groupId	artifactId	version
	org.apache.solr	solr-uima	4.4.0-cdh5.1.3
	org.apache.solr	solr-velocity	4.4.0-cdh5.1.3
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.1.3
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.1.3
	com.ngdata	hbase-indexer-cli	1.5-cdh5.1.3
	com.ngdata	hbase-indexer-common	1.5-cdh5.1.3
	com.ngdata	hbase-indexer-demo	1.5-cdh5.1.3
	com.ngdata	hbase-indexer-dist	1.5-cdh5.1.3
	com.ngdata	hbase-indexer-engine	1.5-cdh5.1.3
	com.ngdata	hbase-indexer-model	1.5-cdh5.1.3
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.1.3
	com.ngdata	hbase-indexer-mr	1.5-cdh5.1.3
	com.ngdata	hbase-indexer-server	1.5-cdh5.1.3
	com.ngdata	hbase-sep-api	1.5-cdh5.1.3
	com.ngdata	hbase-sep-demo	1.5-cdh5.1.3
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.1.3
	com.ngdata	hbase-sep-tools	1.5-cdh5.1.3

CDH 5.1.2

**Note:**

There is no CDH 5.1.1 release. This skip in the CDH 5.x sequence allows the CDH and Cloudera Manager components of Cloudera Enterprise 5.1.2 to have consistent numbering.

The following table lists the project name, groupId, artifactId, and version required to access each CDH 5 artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.1.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.1.2
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.1.2

Project	groupId	artifactId	version
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.3.0-mr1-cdh5.1.2
	org.apache.hadoop	hadoop-core	2.3.0-mr1-cdh5.1.2
	org.apache.hadoop	hadoop-examples	2.3.0-mr1-cdh5.1.2
	org.apache.hadoop	hadoop-minicluster	2.3.0-mr1-cdh5.1.2
	org.apache.hadoop	hadoop-streaming	2.3.0-mr1-cdh5.1.2
	org.apache.hadoop	hadoop-test	2.3.0-mr1-cdh5.1.2
	org.apache.hadoop	hadoop-tools	2.3.0-mr1-cdh5.1.2
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.1.2
	org.apache.hive	hive-beeline	0.12.0-cdh5.1.2
	org.apache.hive	hive-cli	0.12.0-cdh5.1.2
	org.apache.hive	hive-common	0.12.0-cdh5.1.2
	org.apache.hive	hive-contrib	0.12.0-cdh5.1.2
	org.apache.hive	hive-exec	0.12.0-cdh5.1.2
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.1.2
	org.apache.hive	hive-hwi	0.12.0-cdh5.1.2
	org.apache.hive	hive-jdbc	0.12.0-cdh5.1.2
	org.apache.hive	hive-metastore	0.12.0-cdh5.1.2
	org.apache.hive	hive-serde	0.12.0-cdh5.1.2
	org.apache.hive	hive-service	0.12.0-cdh5.1.2
	org.apache.hive	hive-shims	0.12.0-cdh5.1.2
	org.apache.hive	hive-testutils	0.12.0-cdh5.1.2
	org.apache.hivecatalog	hive-hcatalog-core	0.12.0-cdh5.1.2
	org.apache.hivecatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.1.2
	org.apache.hivecatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.1.2
	org.apache.hivecatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.1.2
	org.apache.hivecatalog	hive-webhcat	0.12.0-cdh5.1.2
	org.apache.hivecatalog	hive-webhcat-java-client	0.12.0-cdh5.1.2
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.1.2
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.1.2
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.1.2

Version and Download Information

Project	groupId	artifactId	version
Apache HBase	org.apache.hbase	hbase-client	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-common	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-examples	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-hadoop-compat	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-hadoop2-compat	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-it	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-prefix-tree	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-protocol	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-server	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-shell	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-testing-util	0.98.1-cdh5.1.2
	org.apache.hbase	hbase-thrift	0.98.1-cdh5.1.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.1.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.4-cdh5.1.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.1.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.1.2
	org.apache.pig	pigunit	0.12.0-cdh5.1.2
Apache Flume	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.1.2
	org.apache.flume	flume-ng-core	1.5.0-cdh5.1.2
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.1.2
	org.apache.flume	flume-ng-node	1.5.0-cdh5.1.2
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.1.2
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.1.2
	org.apache.flume	flume-tools	1.5.0-cdh5.1.2
	org.apache.flume	flume-file-channel	1.5.0-cdh5.1.2
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.1.2
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.1.2
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.1.2
	org.apache.flume	flume-avro-source	1.5.0-cdh5.1.2
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.1.2

Project	groupId	artifactId	version
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.1.2
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.1.2
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.1.2
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.1.2
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.1.2
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.1.2
	org.apache.flume	flume-jms-source	1.5.0-cdh5.1.2
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.1.2
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.1.2
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-core	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-examples	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-tools	4.0.0-cdh5.1.2
	org.apache.oozie	oozie-mini	4.0.0-cdh5.1.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.1.2
	org.apache.mahout	mahout-core	0.9-cdh5.1.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.mahout	mahout-examples	0.9-cdh5.1.2
	org.apache.mahout	mahout-integration	0.9-cdh5.1.2
	org.apache.mahout	mahout-math	0.9-cdh5.1.2
	org.apache.mahout	mahout-math-scala	0.9-cdh5.1.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.1.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.1.2
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.1.2
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.1.2
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.1.2
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.1.2
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.1.2
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.1.2

Project	groupId	artifactId	version
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.1.2
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.1.2
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.1.2
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.1.2
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.1.2
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.1.2
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-binding-solr	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-core-common	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-core-model-db	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-core-model-search	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-dist	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-policy-common	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-policy-db	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-policy-search	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-provider-cache	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-provider-common	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-provider-db	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-provider-file	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-tests-hive	1.3.0-cdh5.1.2
	org.apache.sentry	sentry-tests-solr	1.3.0-cdh5.1.2
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.1.2
	com.twitter	parquet-cascading	1.2.5-cdh5.1.2
	com.twitter	parquet-column	1.2.5-cdh5.1.2
	com.twitter	parquet-common	1.2.5-cdh5.1.2
	com.twitter	parquet-encoding	1.2.5-cdh5.1.2
	com.twitter	parquet-format	1.0.0-cdh5.1.2
	com.twitter	parquet-generator	1.2.5-cdh5.1.2
	com.twitter	parquet-hadoop	1.2.5-cdh5.1.2
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.1.2

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-pig	1.2.5-cdh5.1.2
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.1.2
	com.twitter	parquet-scrooge	1.2.5-cdh5.1.2
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.1.2
	com.twitter	parquet-thrift	1.2.5-cdh5.1.2
Llama	com.doudera.llama	llama	1.0.0-cdh5.1.2
Apache Spark	org.apache.spark	spark-bagel_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-catalyst_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-core_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-examples_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-graphx_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-hive_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-mllib_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-repl_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-sql_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-streaming-flume_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-streaming-kafka_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-streaming-mqtt_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-streaming-twitter_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-streaming-zeromq_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-streaming_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-tools_2.10	1.0.0-cdh5.1.2
	org.apache.spark	spark-yarn_2.10	1.0.0-cdh5.1.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.10.0-cdh5.1.2
	org.apache.crunch	crunch-contrib	0.10.0-cdh5.1.2
	org.apache.crunch	crunch-core	0.10.0-cdh5.1.2
	org.apache.crunch	crunch-examples	0.10.0-cdh5.1.2
	org.apache.crunch	crunch-hbase	0.10.0-cdh5.1.2
	org.apache.crunch	crunch-scrunch	0.10.0-cdh5.1.2
	org.apache.crunch	crunch-spark	0.10.0-cdh5.1.2

Project	groupId	artifactId	version
	org.apache.crunch	crunch-test	0.10.0-cdh5.1.2
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.1.2
	org.apache.avro	avro-compiler	1.7.5-cdh5.1.2
	org.apache.avro	avro-ipc	1.7.5-cdh5.1.2
	org.apache.avro	avro-mapred	1.7.5-cdh5.1.2
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.1.2
	org.apache.avro	avro-protobuf	1.7.5-cdh5.1.2
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.1.2
	org.apache.avro	avro-thrift	1.7.5-cdh5.1.2
	org.apache.avro	avro-tools	1.7.5-cdh5.1.2
	org.apache.avro	trevni-avro	1.7.5-cdh5.1.2
	org.apache.avro	trevni-core	1.7.5-cdh5.1.2
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.1.2
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.1.2
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.1.2
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.1.2
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.1.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.1.2
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.1.2
	org.kitesdk	kite-tools	0.10.0-cdh5.1.2
Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-classification	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-core	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-demo	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-facet	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-join	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-memory	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-misc	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-queries	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.1.2
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.1.2

Project	groupId	artifactId	version
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.1.2
	org.apache.solr	solr-cell	4.4.0-cdh5.1.2
	org.apache.solr	solr-clustering	4.4.0-cdh5.1.2
	org.apache.solr	solr-core	4.4.0-cdh5.1.2
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.1.2
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.1.2
	org.apache.solr	solr-langid	4.4.0-cdh5.1.2
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.1.2
	org.apache.solr	solr-solrj	4.4.0-cdh5.1.2
	org.apache.solr	solr-test-framework	4.4.0-cdh5.1.2
	org.apache.solr	solr-uima	4.4.0-cdh5.1.2
	org.apache.solr	solr-velocity	4.4.0-cdh5.1.2
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.1.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.1.2
	com.ngdata	hbase-indexer-cli	1.5-cdh5.1.2
	com.ngdata	hbase-indexer-common	1.5-cdh5.1.2
	com.ngdata	hbase-indexer-demo	1.5-cdh5.1.2
	com.ngdata	hbase-indexer-dist	1.5-cdh5.1.2
	com.ngdata	hbase-indexer-engine	1.5-cdh5.1.2
	com.ngdata	hbase-indexer-model	1.5-cdh5.1.2
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.1.2
	com.ngdata	hbase-indexer-mr	1.5-cdh5.1.2
	com.ngdata	hbase-indexer-server	1.5-cdh5.1.2
	com.ngdata	hbase-sep-api	1.5-cdh5.1.2
	com.ngdata	hbase-sep-demo	1.5-cdh5.1.2
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.1.2
	com.ngdata	hbase-sep-tools	1.5-cdh5.1.2

CDH 5.1.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH 5 artifact.

Version and Download Information

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.1.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.1.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-core	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-examples	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-minicluster	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-streaming	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-test	2.3.0-cdh5.1.0
	org.apache.hadoop	hadoop-tools	2.3.0-cdh5.1.0
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.1.0
	org.apache.hive	hive-beeline	0.12.0-cdh5.1.0
	org.apache.hive	hive-cli	0.12.0-cdh5.1.0
	org.apache.hive	hive-common	0.12.0-cdh5.1.0
	org.apache.hive	hive-contrib	0.12.0-cdh5.1.0
	org.apache.hive	hive-exec	0.12.0-cdh5.1.0
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.1.0
	org.apache.hive	hive-hwi	0.12.0-cdh5.1.0
	org.apache.hive	hive-jdbc	0.12.0-cdh5.1.0
	org.apache.hive	hive-metastore	0.12.0-cdh5.1.0
	org.apache.hive	hive-serde	0.12.0-cdh5.1.0
	org.apache.hive	hive-service	0.12.0-cdh5.1.0
	org.apache.hive	hive-shims	0.12.0-cdh5.1.0
	org.apache.hive	hive-testutils	0.12.0-cdh5.1.0
	org.apache.hive	hive-hcatalog-core	0.12.0-cdh5.1.0
	org.apache.hive	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.1.0
	org.apache.hive	hive-hcatalog-pig-adapter	0.12.0-cdh5.1.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.1.0
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.1.0
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.1.0
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.1.0
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.1.0
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.1.0
Apache HBase	org.apache.hbase	hbase-client	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-common	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-examples	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-hadoop-compat	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-hadoop2-compat	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-it	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-prefix-tree	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-protocol	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-server	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-shell	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-testing-util	0.98.1-cdh5.1.0
	org.apache.hbase	hbase-thrift	0.98.1-cdh5.1.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.1.0
Apache Sqoop	org.apache.sqoop	sqoop	1.4.4-cdh5.1.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.1.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.1.0
	org.apache.pig	pigunit	0.12.0-cdh5.1.0
Apache Flume	org.apache.flume	flume-ng-configuration	1.5.0-cdh5.1.0
	org.apache.flume	flume-ng-core	1.5.0-cdh5.1.0
	org.apache.flume	flume-ng-embedded-agent	1.5.0-cdh5.1.0
	org.apache.flume	flume-ng-node	1.5.0-cdh5.1.0
	org.apache.flume	flume-ng-sdk	1.5.0-cdh5.1.0
	org.apache.flume	flume-ng-tests	1.5.0-cdh5.1.0
	org.apache.flume	flume-tools	1.5.0-cdh5.1.0

Project	groupId	artifactId	version
	org.apache.flume	flume-file-channel	1.5.0-cdh5.1.0
	org.apache.flume	flume-jdbc-channel	1.5.0-cdh5.1.0
	org.apache.flume	flume-spillable-memory-channel	1.5.0-cdh5.1.0
	org.apache.flume	flume-ng-log4jappender	1.5.0-cdh5.1.0
	org.apache.flume	flume-avro-source	1.5.0-cdh5.1.0
	org.apache.flume	flume-thrift-source	1.5.0-cdh5.1.0
	org.apache.flume	flume-dataset-sink	1.5.0-cdh5.1.0
	org.apache.flume	flume-hdfs-sink	1.5.0-cdh5.1.0
	org.apache.flume	flume-irc-sink	1.5.0-cdh5.1.0
	org.apache.flume	flume-ng-elasticsearch-sink	1.5.0-cdh5.1.0
	org.apache.flume	flume-ng-hbase-sink	1.5.0-cdh5.1.0
	org.apache.flume	flume-ng-morphline-solr-sink	1.5.0-cdh5.1.0
	org.apache.flume	flume-jms-source	1.5.0-cdh5.1.0
	org.apache.flume	flume-scribe-source	1.5.0-cdh5.1.0
	org.apache.flume	flume-twitter-source	1.5.0-cdh5.1.0
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-core	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-examples	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.1.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.1.0
	org.apache.oozie	oozie-tools	4.0.0-cdh5.1.0
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.1.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.9-cdh5.1.0
	org.apache.mahout	mahout-core	0.9-cdh5.1.0
	org.apache.mahout	mahout-examples	0.9-cdh5.1.0
	org.apache.mahout	mahout-integration	0.9-cdh5.1.0
	org.apache.mahout	mahout-math	0.9-cdh5.1.0
	org.apache.mahout	mahout-math-scala	0.9-cdh5.1.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.1.0
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.1.0

Project	groupId	artifactId	version	
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.1.0	
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.1.0	
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.1.0	
	Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.3.0-cdh5.1.0
		org.apache.sentry	sentry-binding-solr	1.3.0-cdh5.1.0
org.apache.sentry		sentry-core-common	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-core-model-db	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-core-model-search	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-dist	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-policy-common	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-policy-db	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-policy-search	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-provider-cache	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-provider-common	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-provider-db	1.3.0-cdh5.1.0	
org.apache.sentry		sentry-provider-file	1.3.0-cdh5.1.0	
org.apache.sentry	sentry-tests-hive	1.3.0-cdh5.1.0		
org.apache.sentry	sentry-tests-solr	1.3.0-cdh5.1.0		
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.1.0	
	com.twitter	parquet-cascading	1.2.5-cdh5.1.0	
	com.twitter	parquet-column	1.2.5-cdh5.1.0	

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-common	1.2.5-cdh5.1.0
	com.twitter	parquet-encoding	1.2.5-cdh5.1.0
	com.twitter	parquet-format	1.0.0-cdh5.1.0
	com.twitter	parquet-generator	1.2.5-cdh5.1.0
	com.twitter	parquet-hadoop	1.2.5-cdh5.1.0
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.1.0
	com.twitter	parquet-pig	1.2.5-cdh5.1.0
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.1.0
	com.twitter	parquet-scrooge	1.2.5-cdh5.1.0
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.1.0
	com.twitter	parquet-thrift	1.2.5-cdh5.1.0
Llama	com.doudera.llama	llama	1.0.0-cdh5.1.0
Apache Spark	org.apache.spark	spark-bagel_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-catalyst_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-core_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-examples_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-graphx_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-hive_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-mllib_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-repl_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-sql_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-streaming-flume_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-streaming-kafka_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-streaming-mqtt_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-streaming-twitter_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-streaming-zeromq_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-streaming_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-tools_2.10	1.0.0-cdh5.1.0
	org.apache.spark	spark-yarn_2.10	1.0.0-cdh5.1.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.10.0-cdh5.1.0

Project	groupId	artifactId	version
	org.apache.crunch	crunch-contrib	0.10.0-cdh5.1.0
	org.apache.crunch	crunch-core	0.10.0-cdh5.1.0
	org.apache.crunch	crunch-examples	0.10.0-cdh5.1.0
	org.apache.crunch	crunch-hbase	0.10.0-cdh5.1.0
	org.apache.crunch	crunch-scrunch	0.10.0-cdh5.1.0
	org.apache.crunch	crunch-spark	0.10.0-cdh5.1.0
	org.apache.crunch	crunch-test	0.10.0-cdh5.1.0
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.1.0
	org.apache.avro	avro-compiler	1.7.5-cdh5.1.0
	org.apache.avro	avro-ipc	1.7.5-cdh5.1.0
	org.apache.avro	avro-mapred	1.7.5-cdh5.1.0
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.1.0
	org.apache.avro	avro-protobuf	1.7.5-cdh5.1.0
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.1.0
	org.apache.avro	avro-thrift	1.7.5-cdh5.1.0
	org.apache.avro	avro-tools	1.7.5-cdh5.1.0
	org.apache.avro	trevni-avro	1.7.5-cdh5.1.0
	org.apache.avro	trevni-core	1.7.5-cdh5.1.0
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.1.0
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.1.0
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.1.0
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.1.0
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.1.0

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.1.0
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.1.0
	org.kitesdk	kite-tools	0.10.0-cdh5.1.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-classification	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-core	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-demo	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-facet	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-join	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-memory	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-misc	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-queries	4.4.0-cdh5.1.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.1.0
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.1.0
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.1.0
	org.apache.solr	solr-cell	4.4.0-cdh5.1.0
	org.apache.solr	solr-clustering	4.4.0-cdh5.1.0
	org.apache.solr	solr-core	4.4.0-cdh5.1.0
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.1.0
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.1.0
	org.apache.solr	solr-langid	4.4.0-cdh5.1.0
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.1.0
	org.apache.solr	solr-solrj	4.4.0-cdh5.1.0
	org.apache.solr	solr-test-framework	4.4.0-cdh5.1.0
	org.apache.solr	solr-uima	4.4.0-cdh5.1.0
	org.apache.solr	solr-velocity	4.4.0-cdh5.1.0
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.1.0
HBase Indexer	com.ngdata	hbase-indexer-all	1.5-cdh5.1.0
	com.ngdata	hbase-indexer-cli	1.5-cdh5.1.0
	com.ngdata	hbase-indexer-common	1.5-cdh5.1.0
	com.ngdata	hbase-indexer-demo	1.5-cdh5.1.0
	com.ngdata	hbase-indexer-dist	1.5-cdh5.1.0
	com.ngdata	hbase-indexer-engine	1.5-cdh5.1.0
	com.ngdata	hbase-indexer-model	1.5-cdh5.1.0
	com.ngdata	hbase-indexer-morphlines	1.5-cdh5.1.0
	com.ngdata	hbase-indexer-mr	1.5-cdh5.1.0
	com.ngdata	hbase-indexer-server	1.5-cdh5.1.0
	com.ngdata	hbase-sep-api	1.5-cdh5.1.0

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-demo	1.5-cdh5.1.0
	com.ngdata	hbase-sep-impl-common	1.5-cdh5.1.0
	com.ngdata	hbase-sep-tools	1.5-cdh5.1.0

Maven Artifacts for CDH 5.0.x Releases CDH 5.0.6

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.0.6

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.0.6
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.0.6
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	230-mr1-cdh5.0.6
	org.apache.hadoop	hadoop-core	230-mr1-cdh5.0.6
	org.apache.hadoop	hadoop-examples	230-mr1-cdh5.0.6
	org.apache.hadoop	hadoop-minicluster	230-mr1-cdh5.0.6
	org.apache.hadoop	hadoop-streaming	230-mr1-cdh5.0.6
	org.apache.hadoop	hadoop-test	230-mr1-cdh5.0.6
	org.apache.hadoop	hadoop-tools	230-mr1-cdh5.0.6
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.0.6
	org.apache.hive	hive-beeline	0.12.0-cdh5.0.6
	org.apache.hive	hive-cli	0.12.0-cdh5.0.6
	org.apache.hive	hive-common	0.12.0-cdh5.0.6
	org.apache.hive	hive-contrib	0.12.0-cdh5.0.6
	org.apache.hive	hive-exec	0.12.0-cdh5.0.6
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.0.6
	org.apache.hive	hive-hwi	0.12.0-cdh5.0.6
	org.apache.hive	hive-jdbc	0.12.0-cdh5.0.6
	org.apache.hive	hive-metastore	0.12.0-cdh5.0.6
	org.apache.hive	hive-serde	0.12.0-cdh5.0.6

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-service	0.12.0-cdh5.0.6
	org.apache.hive	hive-shims	0.12.0-cdh5.0.6
	org.apache.hive	hive-testutils	0.12.0-cdh5.0.6
	org.apache.hcatalog	hive-hcatalog-core	0.12.0-cdh5.0.6
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.0.6
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.0.6
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.0.6
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.0.6
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.0.6
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.0.6
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.0.6
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.0.6
Apache HBase	org.apache.hbase	hbase-client	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-common	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-examples	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-hadoop-compat	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-hadoop2-compat	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-it	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-prefix-tree	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-protocol	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-server	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-shell	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-testing-util	0.96.1.1-cdh5.0.6
	org.apache.hbase	hbase-thrift	0.96.1.1-cdh5.0.6
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.0.6
Apache Sqoop	org.apache.sqoop	sqoop	1.4.4-cdh5.0.6
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.0.6
	org.apache.pig	pigsmoke	0.12.0-cdh5.0.6
	org.apache.pig	pigunit	0.12.0-cdh5.0.6

Project	groupId	artifactId	version
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.4.0-cdh5.0.6
	org.apache.flume	flume-ng-core	1.4.0-cdh5.0.6
	org.apache.flume	flume-ng-embedded-agent	1.4.0-cdh5.0.6
	org.apache.flume	flume-ng-node	1.4.0-cdh5.0.6
	org.apache.flume	flume-ng-sdk	1.4.0-cdh5.0.6
	org.apache.flume	flume-ng-tests	1.4.0-cdh5.0.6
	org.apache.flume	flume-tools	1.4.0-cdh5.0.6
	org.apache.flume	flume-file-channel	1.4.0-cdh5.0.6
	org.apache.flume	flume-jdbc-channel	1.4.0-cdh5.0.6
	org.apache.flume	flume-ng-log4jappender	1.4.0-cdh5.0.6
	org.apache.flume	flume-avro-source	1.4.0-cdh5.0.6
	org.apache.flume	flume-thrift-source	1.4.0-cdh5.0.6
	org.apache.flume	flume-dataset-sink	1.4.0-cdh5.0.6
	org.apache.flume	flume-hdfs-sink	1.4.0-cdh5.0.6
	org.apache.flume	flume-irc-sink	1.4.0-cdh5.0.6
	org.apache.flume	flume-ng-elasticsearch-sink	1.4.0-cdh5.0.6
	org.apache.flume	flume-ng-hbase-sink	1.4.0-cdh5.0.6
	org.apache.flume	flume-ng-morphline-solr-sink	1.4.0-cdh5.0.6
	org.apache.flume	flume-jms-source	1.4.0-cdh5.0.6
	org.apache.flume	flume-scribe-source	1.4.0-cdh5.0.6
org.apache.flume	flume-twitter-source	1.4.0-cdh5.0.6	
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-core	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-examples	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.0.6
org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.0.6	

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.0.6
	org.apache.oozie	oozie-tools	4.0.0-cdh5.0.6
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.0.6
Apache Mahout	org.apache.mahout	mahout-buildtools	0.8-cdh5.0.6
	org.apache.mahout	mahout-core	0.8-cdh5.0.6
	org.apache.mahout	mahout-examples	0.8-cdh5.0.6
	org.apache.mahout	mahout-integration	0.8-cdh5.0.6
	org.apache.mahout	mahout-math	0.8-cdh5.0.6
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-chef	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-cli	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-core	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-druid	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-examples	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-hama	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.0.6

Project	groupId	artifactId	version
	org.apache.whirr	whirr-pig	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-solr	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.0.6
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.0.6
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.0.6
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.0.6
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.0.6
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-binding-solr	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-core-common	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-core-model-db	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-core-model-search	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-dist	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-policy-common	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-policy-db	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-policy-search	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-provider-common	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-provider-file	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-tests-hive	1.2.0-cdh5.0.6
	org.apache.sentry	sentry-tests-solr	1.2.0-cdh5.0.6

Version and Download Information

Project	groupId	artifactId	version
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.0.6
	com.twitter	parquet-cascading	1.2.5-cdh5.0.6
	com.twitter	parquet-column	1.2.5-cdh5.0.6
	com.twitter	parquet-common	1.2.5-cdh5.0.6
	com.twitter	parquet-encoding	1.2.5-cdh5.0.6
	com.twitter	parquet-format	1.0.0-cdh5.0.6
	com.twitter	parquet-generator	1.2.5-cdh5.0.6
	com.twitter	parquet-hadoop	1.2.5-cdh5.0.6
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.0.6
	com.twitter	parquet-pig	1.2.5-cdh5.0.6
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.0.6
	com.twitter	parquet-scrooge	1.2.5-cdh5.0.6
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.0.6
	com.twitter	parquet-thrift	1.2.5-cdh5.0.6
	Llama	com.douderallama	llama
Apache Spark	org.apache.spark	spark-bagel_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-core_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-examples_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-graphx_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-mllib_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-repl_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-streaming-flume_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-streaming-kafka_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-streaming-mqtt_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-streaming-twitter_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-streaming-zeromq_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-streaming_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-tools_2.10	0.9.0-cdh5.0.6
	org.apache.spark	spark-yarn_2.10	0.9.0-cdh5.0.6
	Apache Crunch	org.apache.crunch	crunch-archetype

Project	groupId	artifactId	version
	org.apache.crunch	crunch-contrib	0.9.0-cdh5.0.6
	org.apache.crunch	crunch-core	0.9.0-cdh5.0.6
	org.apache.crunch	crunch-examples	0.9.0-cdh5.0.6
	org.apache.crunch	crunch-hbase	0.9.0-cdh5.0.6
	org.apache.crunch	crunch-scrunch	0.9.0-cdh5.0.6
	org.apache.crunch	crunch-spark	0.9.0-cdh5.0.6
	org.apache.crunch	crunch-test	0.9.0-cdh5.0.6
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.0.6
	org.apache.avro	avro-compiler	1.7.5-cdh5.0.6
	org.apache.avro	avro-ipc	1.7.5-cdh5.0.6
	org.apache.avro	avro-mapred	1.7.5-cdh5.0.6
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.0.6
	org.apache.avro	avro-protobuf	1.7.5-cdh5.0.6
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.0.6
	org.apache.avro	avro-thrift	1.7.5-cdh5.0.6
	org.apache.avro	avro-tools	1.7.5-cdh5.0.6
	org.apache.avro	trevni-avro	1.7.5-cdh5.0.6
	org.apache.avro	trevni-core	1.7.5-cdh5.0.6
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.0.6
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.0.6
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.0.6
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.0.6
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.0.6

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.0.6
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.0.6
	org.kitesdk	kite-tools	0.10.0-cdh5.0.6
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-classification	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-core	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-demo	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-facet	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-join	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-memory	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-misc	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-queries	4.4.0-cdh5.0.6

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.0.6
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.0.6
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.0.6
	org.apache.solr	solr-cell	4.4.0-cdh5.0.6
	org.apache.solr	solr-clustering	4.4.0-cdh5.0.6
	org.apache.solr	solr-core	4.4.0-cdh5.0.6
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.0.6
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.0.6
	org.apache.solr	solr-langid	4.4.0-cdh5.0.6
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.0.6
	org.apache.solr	solr-solrj	4.4.0-cdh5.0.6
	org.apache.solr	solr-test-framework	4.4.0-cdh5.0.6
	org.apache.solr	solr-uima	4.4.0-cdh5.0.6
	org.apache.solr	solr-velocity	4.4.0-cdh5.0.6
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.0.6
HBase Indexer	com.ngdata	hbase-indexer-all	1.3-cdh5.0.6
	com.ngdata	hbase-indexer-cli	1.3-cdh5.0.6
	com.ngdata	hbase-indexer-common	1.3-cdh5.0.6
	com.ngdata	hbase-indexer-demo	1.3-cdh5.0.6
	com.ngdata	hbase-indexer-dist	1.3-cdh5.0.6
	com.ngdata	hbase-indexer-engine	1.3-cdh5.0.6
	com.ngdata	hbase-indexer-model	1.3-cdh5.0.6
	com.ngdata	hbase-indexer-morphlines	1.3-cdh5.0.6
	com.ngdata	hbase-indexer-mr	1.3-cdh5.0.6
	com.ngdata	hbase-indexer-server	1.3-cdh5.0.6
	com.ngdata	hbase-sep-api	1.3-cdh5.0.6

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-demo	1.3-cdh5.0.6
	com.ngdata	hbase-sep-impl-common	1.3-cdh5.0.6
	com.ngdata	hbase-sep-tools	1.3-cdh5.0.6

CDH 5.0.5

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.0.5

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.0.5
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.0.5
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	230-mr1-cdh5.0.5
	org.apache.hadoop	hadoop-core	230-mr1-cdh5.0.5
	org.apache.hadoop	hadoop-examples	230-mr1-cdh5.0.5
	org.apache.hadoop	hadoop-minicluster	230-mr1-cdh5.0.5
	org.apache.hadoop	hadoop-streaming	230-mr1-cdh5.0.5
	org.apache.hadoop	hadoop-test	230-mr1-cdh5.0.5
	org.apache.hadoop	hadoop-tools	230-mr1-cdh5.0.5
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.0.5
	org.apache.hive	hive-beeline	0.12.0-cdh5.0.5
	org.apache.hive	hive-cli	0.12.0-cdh5.0.5
	org.apache.hive	hive-common	0.12.0-cdh5.0.5
	org.apache.hive	hive-contrib	0.12.0-cdh5.0.5
	org.apache.hive	hive-exec	0.12.0-cdh5.0.5
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.0.5
	org.apache.hive	hive-hwi	0.12.0-cdh5.0.5
	org.apache.hive	hive-jdbc	0.12.0-cdh5.0.5
	org.apache.hive	hive-metastore	0.12.0-cdh5.0.5
	org.apache.hive	hive-serde	0.12.0-cdh5.0.5

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-service	0.12.0-cdh5.0.5
	org.apache.hive	hive-shims	0.12.0-cdh5.0.5
	org.apache.hive	hive-testutils	0.12.0-cdh5.0.5
	org.apache.hcatalog	hive-hcatalog-core	0.12.0-cdh5.0.5
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.0.5
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.0.5
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.0.5
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.0.5
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.0.5
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.0.5
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.0.5
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.0.5
Apache HBase	org.apache.hbase	hbase-client	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-common	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-examples	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-hadoop-compat	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-hadoop2-compat	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-it	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-prefix-tree	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-protocol	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-server	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-shell	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-testing-util	0.96.1.1-cdh5.0.5
	org.apache.hbase	hbase-thrift	0.96.1.1-cdh5.0.5
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.0.5
Apache Sqoop	org.apache.sqoop	sqoop	1.4.4-cdh5.0.5
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.0.5
	org.apache.pig	pigsmoke	0.12.0-cdh5.0.5
	org.apache.pig	pigunit	0.12.0-cdh5.0.5

Project	groupId	artifactId	version
Apache Flume 1.x	org.apache.flume	flume-ng-configuration	1.4.0-cdh5.0.5
	org.apache.flume	flume-ng-core	1.4.0-cdh5.0.5
	org.apache.flume	flume-ng-embedded-agent	1.4.0-cdh5.0.5
	org.apache.flume	flume-ng-node	1.4.0-cdh5.0.5
	org.apache.flume	flume-ng-sdk	1.4.0-cdh5.0.5
	org.apache.flume	flume-ng-tests	1.4.0-cdh5.0.5
	org.apache.flume	flume-tools	1.4.0-cdh5.0.5
	org.apache.flume	flume-file-channel	1.4.0-cdh5.0.5
	org.apache.flume	flume-jdbc-channel	1.4.0-cdh5.0.5
	org.apache.flume	flume-ng-log4jappender	1.4.0-cdh5.0.5
	org.apache.flume	flume-avro-source	1.4.0-cdh5.0.5
	org.apache.flume	flume-thrift-source	1.4.0-cdh5.0.5
	org.apache.flume	flume-dataset-sink	1.4.0-cdh5.0.5
	org.apache.flume	flume-hdfs-sink	1.4.0-cdh5.0.5
	org.apache.flume	flume-irc-sink	1.4.0-cdh5.0.5
	org.apache.flume	flume-ng-elasticsearch-sink	1.4.0-cdh5.0.5
	org.apache.flume	flume-ng-hbase-sink	1.4.0-cdh5.0.5
	org.apache.flume	flume-ng-morphline-solr-sink	1.4.0-cdh5.0.5
	org.apache.flume	flume-jms-source	1.4.0-cdh5.0.5
	org.apache.flume	flume-scribe-source	1.4.0-cdh5.0.5
org.apache.flume	flume-twitter-source	1.4.0-cdh5.0.5	
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-core	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-examples	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.0.5
org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.0.5	

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.0.5
	org.apache.oozie	oozie-tools	4.0.0-cdh5.0.5
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.0.5
Apache Mahout	org.apache.mahout	mahout-buildtools	0.8-cdh5.0.5
	org.apache.mahout	mahout-core	0.8-cdh5.0.5
	org.apache.mahout	mahout-examples	0.8-cdh5.0.5
	org.apache.mahout	mahout-integration	0.8-cdh5.0.5
	org.apache.mahout	mahout-math	0.8-cdh5.0.5
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-chef	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-cli	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-core	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-druid	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-examples	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-hama	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.0.5

Project	groupId	artifactId	version
	org.apache.whirr	whirr-pig	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-solr	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.0.5
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.0.5
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.0.5
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.0.5
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.0.5
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-binding-solr	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-core-common	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-core-model-db	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-core-model-search	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-dist	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-policy-common	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-policy-db	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-policy-search	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-provider-common	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-provider-file	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-tests-hive	1.2.0-cdh5.0.5
	org.apache.sentry	sentry-tests-solr	1.2.0-cdh5.0.5

Version and Download Information

Project	groupId	artifactId	version
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.0.5
	com.twitter	parquet-cascading	1.2.5-cdh5.0.5
	com.twitter	parquet-column	1.2.5-cdh5.0.5
	com.twitter	parquet-common	1.2.5-cdh5.0.5
	com.twitter	parquet-encoding	1.2.5-cdh5.0.5
	com.twitter	parquet-format	1.0.0-cdh5.0.5
	com.twitter	parquet-generator	1.2.5-cdh5.0.5
	com.twitter	parquet-hadoop	1.2.5-cdh5.0.5
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.0.5
	com.twitter	parquet-pig	1.2.5-cdh5.0.5
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.0.5
	com.twitter	parquet-scrooge	1.2.5-cdh5.0.5
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.0.5
	com.twitter	parquet-thrift	1.2.5-cdh5.0.5
Llama	com.douderallama	llama	1.0.0-cdh5.0.5
Apache Spark	org.apache.spark	spark-bagel_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-core_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-examples_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-graphx_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-mllib_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-repl_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-streaming-flume_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-streaming-kafka_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-streaming-mqtt_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-streaming-twitter_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-streaming-zeromq_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-streaming_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-tools_2.10	0.9.0-cdh5.0.5
	org.apache.spark	spark-yarn_2.10	0.9.0-cdh5.0.5
Apache Crunch	org.apache.crunch	crunch-archetype	0.9.0-cdh5.0.5

Project	groupId	artifactId	version
	org.apache.crunch	crunch-contrib	0.9.0-cdh5.0.5
	org.apache.crunch	crunch-core	0.9.0-cdh5.0.5
	org.apache.crunch	crunch-examples	0.9.0-cdh5.0.5
	org.apache.crunch	crunch-hbase	0.9.0-cdh5.0.5
	org.apache.crunch	crunch-scrunch	0.9.0-cdh5.0.5
	org.apache.crunch	crunch-spark	0.9.0-cdh5.0.5
	org.apache.crunch	crunch-test	0.9.0-cdh5.0.5
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.0.5
	org.apache.avro	avro-compiler	1.7.5-cdh5.0.5
	org.apache.avro	avro-ipc	1.7.5-cdh5.0.5
	org.apache.avro	avro-mapred	1.7.5-cdh5.0.5
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.0.5
	org.apache.avro	avro-protobuf	1.7.5-cdh5.0.5
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.0.5
	org.apache.avro	avro-thrift	1.7.5-cdh5.0.5
	org.apache.avro	avro-tools	1.7.5-cdh5.0.5
	org.apache.avro	trevni-avro	1.7.5-cdh5.0.5
	org.apache.avro	trevni-core	1.7.5-cdh5.0.5
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.0.5
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.0.5
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.0.5
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.0.5
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.0.5

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.0.5
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.0.5
	org.kitesdk	kite-tools	0.10.0-cdh5.0.5
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-classification	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-core	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-demo	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-facet	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-join	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-memory	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-misc	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-queries	4.4.0-cdh5.0.5

Project	groupId	artifactId	version
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.0.5
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.0.5
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.0.5
	org.apache.solr	solr-cell	4.4.0-cdh5.0.5
	org.apache.solr	solr-clustering	4.4.0-cdh5.0.5
	org.apache.solr	solr-core	4.4.0-cdh5.0.5
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.0.5
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.0.5
	org.apache.solr	solr-langid	4.4.0-cdh5.0.5
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.0.5
	org.apache.solr	solr-solrj	4.4.0-cdh5.0.5
	org.apache.solr	solr-test-framework	4.4.0-cdh5.0.5
	org.apache.solr	solr-uima	4.4.0-cdh5.0.5
	org.apache.solr	solr-velocity	4.4.0-cdh5.0.5
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.0.5
HBase Indexer	com.ngdata	hbase-indexer-all	1.3-cdh5.0.5
	com.ngdata	hbase-indexer-cli	1.3-cdh5.0.5
	com.ngdata	hbase-indexer-common	1.3-cdh5.0.5
	com.ngdata	hbase-indexer-demo	1.3-cdh5.0.5
	com.ngdata	hbase-indexer-dist	1.3-cdh5.0.5
	com.ngdata	hbase-indexer-engine	1.3-cdh5.0.5
	com.ngdata	hbase-indexer-model	1.3-cdh5.0.5
	com.ngdata	hbase-indexer-morphlines	1.3-cdh5.0.5
	com.ngdata	hbase-indexer-mr	1.3-cdh5.0.5
	com.ngdata	hbase-indexer-server	1.3-cdh5.0.5
	com.ngdata	hbase-sep-api	1.3-cdh5.0.5

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-demo	1.3-cdh5.0.5
	com.ngdata	hbase-sep-impl-common	1.3-cdh5.0.5
	com.ngdata	hbase-sep-tools	1.3-cdh5.0.5

CDH 5.0.4

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.0.4

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.0.4
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.0.4
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	230-mr1-cdh5.0.4
	org.apache.hadoop	hadoop-core	230-mr1-cdh5.0.4
	org.apache.hadoop	hadoop-examples	230-mr1-cdh5.0.4
	org.apache.hadoop	hadoop-minicluster	230-mr1-cdh5.0.4
	org.apache.hadoop	hadoop-streaming	230-mr1-cdh5.0.4
	org.apache.hadoop	hadoop-test	230-mr1-cdh5.0.4
	org.apache.hadoop	hadoop-tools	230-mr1-cdh5.0.4
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.0.4
	org.apache.hive	hive-beeline	0.12.0-cdh5.0.4
	org.apache.hive	hive-cli	0.12.0-cdh5.0.4
	org.apache.hive	hive-common	0.12.0-cdh5.0.4
	org.apache.hive	hive-contrib	0.12.0-cdh5.0.4
	org.apache.hive	hive-exec	0.12.0-cdh5.0.4
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.0.4
	org.apache.hive	hive-hwi	0.12.0-cdh5.0.4
	org.apache.hive	hive-jdbc	0.12.0-cdh5.0.4
	org.apache.hive	hive-metastore	0.12.0-cdh5.0.4
	org.apache.hive	hive-serde	0.12.0-cdh5.0.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-service	0.12.0-cdh5.0.4
	org.apache.hive	hive-shims	0.12.0-cdh5.0.4
	org.apache.hive	hive-testutils	0.12.0-cdh5.0.4
	org.apache.hcatalog	hive-hcatalog-core	0.12.0-cdh5.0.4
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.0.4
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.0.4
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.0.4
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.0.4
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.0.4
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.0.4
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.0.4
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.0.4
Apache HBase	org.apache.hbase	hbase-client	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-common	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-examples	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-hadoop-compat	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-hadoop2-compat	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-it	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-prefix-tree	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-protocol	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-server	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-shell	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-testing-util	0.96.1.1-cdh5.0.4
	org.apache.hbase	hbase-thrift	0.96.1.1-cdh5.0.4
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.0.4
Sqoop	org.apache.sqoop	sqoop	1.4.4-cdh5.0.4
Pig	org.apache.pig	pig	0.12.0-cdh5.0.4
	org.apache.pig	pigsmoke	0.12.0-cdh5.0.4
	org.apache.pig	pigunit	0.12.0-cdh5.0.4
Apache Flume	org.apache.flume	flume-ng-configuration	1.4.0-cdh5.0.4

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-core	1.4.0-cdh5.0.4
	org.apache.flume	flume-ng-embedded-agent	1.4.0-cdh5.0.4
	org.apache.flume	flume-ng-node	1.4.0-cdh5.0.4
	org.apache.flume	flume-ng-sdk	1.4.0-cdh5.0.4
	org.apache.flume	flume-ng-tests	1.4.0-cdh5.0.4
	org.apache.flume	flume-tools	1.4.0-cdh5.0.4
	org.apache.flume	flume-file-channel	1.4.0-cdh5.0.4
	org.apache.flume	flume-jdbc-channel	1.4.0-cdh5.0.4
	org.apache.flume	flume-ng-log4jappender	1.4.0-cdh5.0.4
	org.apache.flume	flume-avro-source	1.4.0-cdh5.0.4
	org.apache.flume	flume-thrift-source	1.4.0-cdh5.0.4
	org.apache.flume	flume-dataset-sink	1.4.0-cdh5.0.4
	org.apache.flume	flume-hdfs-sink	1.4.0-cdh5.0.4
	org.apache.flume	flume-irc-sink	1.4.0-cdh5.0.4
	org.apache.flume	flume-ng-elasticsearch-sink	1.4.0-cdh5.0.4
	org.apache.flume	flume-ng-hbase-sink	1.4.0-cdh5.0.4
	org.apache.flume	flume-ng-morphline-solr-sink	1.4.0-cdh5.0.4
	org.apache.flume	flume-jms-source	1.4.0-cdh5.0.4
	org.apache.flume	flume-scribe-source	1.4.0-cdh5.0.4
	org.apache.flume	flume-twitter-source	1.4.0-cdh5.0.4
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-core	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-examples	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.0.4

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.0.4
	org.apache.oozie	oozie-tools	4.0.0-cdh5.0.4
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.0.4
Apache Mahout	org.apache.mahout	mahout-buildtools	0.8-cdh5.0.4
	org.apache.mahout	mahout-core	0.8-cdh5.0.4
	org.apache.mahout	mahout-examples	0.8-cdh5.0.4
	org.apache.mahout	mahout-integration	0.8-cdh5.0.4
	org.apache.mahout	mahout-math	0.8-cdh5.0.4
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-chef	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-cli	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-core	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-druid	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-examples	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-hama	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-pig	0.9.0-cdh5.0.4

Project	groupId	artifactId	version
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-solr	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.0.4
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.0.4
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.0.4
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.0.4
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.0.4
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-binding-solr	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-core-common	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-core-model-db	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-core-model-search	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-dist	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-policy-common	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-policy-db	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-policy-search	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-provider-common	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-provider-file	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-tests-hive	1.2.0-cdh5.0.4
	org.apache.sentry	sentry-tests-solr	1.2.0-cdh5.0.4
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.0.4

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-cascading	1.2.5-cdh5.0.4
	com.twitter	parquet-column	1.2.5-cdh5.0.4
	com.twitter	parquet-common	1.2.5-cdh5.0.4
	com.twitter	parquet-encoding	1.2.5-cdh5.0.4
	com.twitter	parquet-format	1.0.0-cdh5.0.4
	com.twitter	parquet-generator	1.2.5-cdh5.0.4
	com.twitter	parquet-hadoop	1.2.5-cdh5.0.4
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.0.4
	com.twitter	parquet-pig	1.2.5-cdh5.0.4
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.0.4
	com.twitter	parquet-scrooge	1.2.5-cdh5.0.4
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.0.4
	com.twitter	parquet-thrift	1.2.5-cdh5.0.4
Llama	com.doudera.llama	llama	1.0.0-cdh5.0.4
Apache Spark	org.apache.spark	spark-bagel_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-core_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-examples_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-graphx_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-mllib_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-repl_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-streaming-flume_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-streaming-kafka_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-streaming-mqtt_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-streaming-twitter_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-streaming-zeromq_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-streaming_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-tools_2.10	0.9.0-cdh5.0.4
	org.apache.spark	spark-yarn_2.10	0.9.0-cdh5.0.4
Apache Crunch	org.apache.crunch	crunch-archetype	0.9.0-cdh5.0.4
	org.apache.crunch	crunch-contrib	0.9.0-cdh5.0.4

Project	groupId	artifactId	version
	org.apache.crunch	crunch-core	0.9.0-cdh5.0.4
	org.apache.crunch	crunch-examples	0.9.0-cdh5.0.4
	org.apache.crunch	crunch-hbase	0.9.0-cdh5.0.4
	org.apache.crunch	crunch-scrunch	0.9.0-cdh5.0.4
	org.apache.crunch	crunch-spark	0.9.0-cdh5.0.4
	org.apache.crunch	crunch-test	0.9.0-cdh5.0.4
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.0.4
	org.apache.avro	avro-compiler	1.7.5-cdh5.0.4
	org.apache.avro	avro-ipc	1.7.5-cdh5.0.4
	org.apache.avro	avro-mapred	1.7.5-cdh5.0.4
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.0.4
	org.apache.avro	avro-protobuf	1.7.5-cdh5.0.4
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.0.4
	org.apache.avro	avro-thrift	1.7.5-cdh5.0.4
	org.apache.avro	avro-tools	1.7.5-cdh5.0.4
	org.apache.avro	trevni-avro	1.7.5-cdh5.0.4
	org.apache.avro	trevni-core	1.7.5-cdh5.0.4
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.0.4
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.0.4
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.0.4
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.0.4
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.0.4

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.0.4
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.0.4
	org.kitesdk	kite-tools	0.10.0-cdh5.0.4
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-classification	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-core	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-demo	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-facet	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-join	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-memory	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-misc	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-queries	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.0.4

Project	groupId	artifactId	version
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.0.4
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.0.4
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.0.4
	org.apache.solr	solr-cell	4.4.0-cdh5.0.4
	org.apache.solr	solr-clustering	4.4.0-cdh5.0.4
	org.apache.solr	solr-core	4.4.0-cdh5.0.4
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.0.4
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.0.4
	org.apache.solr	solr-langid	4.4.0-cdh5.0.4
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.0.4
	org.apache.solr	solr-solrj	4.4.0-cdh5.0.4
	org.apache.solr	solr-test-framework	4.4.0-cdh5.0.4
	org.apache.solr	solr-uima	4.4.0-cdh5.0.4
	org.apache.solr	solr-velocity	4.4.0-cdh5.0.4
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.0.4
HBase Indexer	com.ngdata	hbase-indexer-all	1.3-cdh5.0.4
	com.ngdata	hbase-indexer-cli	1.3-cdh5.0.4
	com.ngdata	hbase-indexer-common	1.3-cdh5.0.4
	com.ngdata	hbase-indexer-demo	1.3-cdh5.0.4
	com.ngdata	hbase-indexer-dist	1.3-cdh5.0.4
	com.ngdata	hbase-indexer-engine	1.3-cdh5.0.4
	com.ngdata	hbase-indexer-model	1.3-cdh5.0.4
	com.ngdata	hbase-indexer-morphlines	1.3-cdh5.0.4
	com.ngdata	hbase-indexer-mr	1.3-cdh5.0.4
	com.ngdata	hbase-indexer-server	1.3-cdh5.0.4
	com.ngdata	hbase-sep-api	1.3-cdh5.0.4
	com.ngdata	hbase-sep-demo	1.3-cdh5.0.4

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-impl-common	1.3-cdh5.0.4
	com.ngdata	hbase-sep-tools	1.3-cdh5.0.4

CDH 5.0.3

The following table lists the project name, groupId, artifactId, and version required to access each CDH 5 artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.0.3

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.0.3
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-core	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-examples	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-minicluster	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-streaming	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-test	2.3.0-cdh5.0.3
	org.apache.hadoop	hadoop-tools	2.3.0-cdh5.0.3
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.0.3
	org.apache.hive	hive-beeline	0.12.0-cdh5.0.3
	org.apache.hive	hive-cli	0.12.0-cdh5.0.3
	org.apache.hive	hive-common	0.12.0-cdh5.0.3
	org.apache.hive	hive-contrib	0.12.0-cdh5.0.3
	org.apache.hive	hive-exec	0.12.0-cdh5.0.3
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.0.3
	org.apache.hive	hive-hwi	0.12.0-cdh5.0.3
	org.apache.hive	hive-jdbc	0.12.0-cdh5.0.3
	org.apache.hive	hive-metastore	0.12.0-cdh5.0.3
	org.apache.hive	hive-serde	0.12.0-cdh5.0.3
	org.apache.hive	hive-service	0.12.0-cdh5.0.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-shims	0.12.0-cdh5.0.3
	org.apache.hive	hive-testutils	0.12.0-cdh5.0.3
	org.apache.hive	hive-hcatalog-core	0.12.0-cdh5.0.3
	org.apache.hive	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.0.3
	org.apache.hive	hive-hcatalog-pig-adapter	0.12.0-cdh5.0.3
	org.apache.hive	hive-hcatalog-server-extensions	0.12.0-cdh5.0.3
	org.apache.hive	hive-webhcat	0.12.0-cdh5.0.3
	org.apache.hive	hive-webhcat-java-client	0.12.0-cdh5.0.3
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.0.3
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.0.3
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.0.3
Apache HBase	org.apache.hbase	hbase-client	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-common	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-examples	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-hadoop-compat	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-hadoop2-compat	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-it	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-prefix-tree	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-protocol	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-server	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-shell	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-testing-util	0.96.1.1-cdh5.0.3
	org.apache.hbase	hbase-thrift	0.96.1.1-cdh5.0.3
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.0.3
Apache Sqoop	org.apache.sqoop	sqoop	1.4.4-cdh5.0.3
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.0.3
	org.apache.pig	pigsmoke	0.12.0-cdh5.0.3
	org.apache.pig	pigunit	0.12.0-cdh5.0.3
Apache Flume	org.apache.flume	flume-ng-configuration	1.4.0-cdh5.0.3
	org.apache.flume	flume-ng-core	1.4.0-cdh5.0.3

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-embedded-agent	1.4.0-cdh5.0.3
	org.apache.flume	flume-ng-node	1.4.0-cdh5.0.3
	org.apache.flume	flume-ng-sdk	1.4.0-cdh5.0.3
	org.apache.flume	flume-ng-tests	1.4.0-cdh5.0.3
	org.apache.flume	flume-tools	1.4.0-cdh5.0.3
	org.apache.flume	flume-file-channel	1.4.0-cdh5.0.3
	org.apache.flume	flume-jdbc-channel	1.4.0-cdh5.0.3
	org.apache.flume	flume-ng-log4jappender	1.4.0-cdh5.0.3
	org.apache.flume	flume-avro-source	1.4.0-cdh5.0.3
	org.apache.flume	flume-thrift-source	1.4.0-cdh5.0.3
	org.apache.flume	flume-dataset-sink	1.4.0-cdh5.0.3
	org.apache.flume	flume-hdfs-sink	1.4.0-cdh5.0.3
	org.apache.flume	flume-irc-sink	1.4.0-cdh5.0.3
	org.apache.flume	flume-ng-elasticsearch-sink	1.4.0-cdh5.0.3
	org.apache.flume	flume-ng-hbase-sink	1.4.0-cdh5.0.3
	org.apache.flume	flume-ng-morphline-solr-sink	1.4.0-cdh5.0.3
	org.apache.flume	flume-jms-source	1.4.0-cdh5.0.3
	org.apache.flume	flume-scribe-source	1.4.0-cdh5.0.3
	org.apache.flume	flume-twitter-source	1.4.0-cdh5.0.3
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-core	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-examples	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-hadoop	20141015-cdh5.0.3
	org.apache.oozie	oozie-hadoop-distcp	20141015-cdh5.0.3
	org.apache.oozie	oozie-hadoop-test	20141015-cdh5.0.3
	org.apache.oozie	oozie-hadoop-utils	20141015-cdh5.0.3
	org.apache.oozie	oozie-hbase	08141015-cdh5.0.3
	org.apache.oozie	oozie-hcatalog	02141015-cdh5.0.3
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.0.3

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.0.3
	org.apache.oozie	oozie-tools	4.0.0-cdh5.0.3
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.0.3
Apache Mahout	org.apache.mahout	mahout-buildtools	0.8-cdh5.0.3
	org.apache.mahout	mahout-core	0.8-cdh5.0.3
	org.apache.mahout	mahout-examples	0.8-cdh5.0.3
	org.apache.mahout	mahout-integration	0.8-cdh5.0.3
	org.apache.mahout	mahout-math	0.8-cdh5.0.3
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-chef	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-cli	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-core	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-druid	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-examples	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-hama	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-pig	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.0.3

Project	groupId	artifactId	version
	org.apache.whirr	whirr-solr	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.0.3
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.0.3
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.0.3
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.0.3
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.0.3
Apache Sentry	org.apache.sentry	sentry-binding-hive	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-binding-solr	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-core-common	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-core-model-db	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-core-model-search	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-dist	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-policy-common	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-policy-db	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-policy-search	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-provider-common	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-provider-file	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-tests-hive	1.2.0-cdh5.0.3
	org.apache.sentry	sentry-tests-solr	1.2.0-cdh5.0.3
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.0.3
	com.twitter	parquet-cascading	1.2.5-cdh5.0.3

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-column	1.2.5-cdh5.0.3
	com.twitter	parquet-common	1.2.5-cdh5.0.3
	com.twitter	parquet-encoding	1.2.5-cdh5.0.3
	com.twitter	parquet-format	1.0.0-cdh5.0.3
	com.twitter	parquet-generator	1.2.5-cdh5.0.3
	com.twitter	parquet-hadoop	1.2.5-cdh5.0.3
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.0.3
	com.twitter	parquet-pig	1.2.5-cdh5.0.3
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.0.3
	com.twitter	parquet-scrooge	1.2.5-cdh5.0.3
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.0.3
	com.twitter	parquet-thrift	1.2.5-cdh5.0.3
Llama	com.doudera.llama	llama	1.0.0-cdh5.0.3
Apache Spark	org.apache.spark	spark-bagel_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-core_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-examples_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-graphx_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-mllib_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-repl_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-streaming-flume_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-streaming-kafka_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-streaming-mqtt_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-streaming-twitter_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-streaming-zeromq_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-streaming_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-tools_2.10	0.9.0-cdh5.0.3
	org.apache.spark	spark-yarn_2.10	0.9.0-cdh5.0.3
Apache Crunch	org.apache.crunch	crunch-archetype	0.9.0-cdh5.0.3
	org.apache.crunch	crunch-contrib	0.9.0-cdh5.0.3
	org.apache.crunch	crunch-core	0.9.0-cdh5.0.3

Project	groupId	artifactId	version
	org.apache.crunch	crunch-examples	0.9.0-cdh5.0.3
	org.apache.crunch	crunch-hbase	0.9.0-cdh5.0.3
	org.apache.crunch	crunch-scrunch	0.9.0-cdh5.0.3
	org.apache.crunch	crunch-spark	0.9.0-cdh5.0.3
	org.apache.crunch	crunch-test	0.9.0-cdh5.0.3
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.0.3
	org.apache.avro	avro-compiler	1.7.5-cdh5.0.3
	org.apache.avro	avro-ipc	1.7.5-cdh5.0.3
	org.apache.avro	avro-mapred	1.7.5-cdh5.0.3
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.0.3
	org.apache.avro	avro-protobuf	1.7.5-cdh5.0.3
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.0.3
	org.apache.avro	avro-thrift	1.7.5-cdh5.0.3
	org.apache.avro	avro-tools	1.7.5-cdh5.0.3
	org.apache.avro	trevni-avro	1.7.5-cdh5.0.3
	org.apache.avro	trevni-core	1.7.5-cdh5.0.3
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.0.3
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.0.3
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.0.3
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.0.3
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.0.3

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.0.3
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.0.3
	org.kitesdk	kite-tools	0.10.0-cdh5.0.3
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-classification	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-core	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-demo	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-facet	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-join	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-memory	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-misc	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-queries	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.0.3

Project	groupId	artifactId	version
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.0.3
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.0.3
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.0.3
	org.apache.solr	solr-cell	4.4.0-cdh5.0.3
	org.apache.solr	solr-clustering	4.4.0-cdh5.0.3
	org.apache.solr	solr-core	4.4.0-cdh5.0.3
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.0.3
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.0.3
	org.apache.solr	solr-langid	4.4.0-cdh5.0.3
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.0.3
	org.apache.solr	solr-solrj	4.4.0-cdh5.0.3
	org.apache.solr	solr-test-framework	4.4.0-cdh5.0.3
	org.apache.solr	solr-uima	4.4.0-cdh5.0.3
	org.apache.solr	solr-velocity	4.4.0-cdh5.0.3
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.0.3
HBase Indexer	com.ngdata	hbase-indexer-all	1.3-cdh5.0.3
	com.ngdata	hbase-indexer-cli	1.3-cdh5.0.3
	com.ngdata	hbase-indexer-common	1.3-cdh5.0.3
	com.ngdata	hbase-indexer-demo	1.3-cdh5.0.3
	com.ngdata	hbase-indexer-dist	1.3-cdh5.0.3
	com.ngdata	hbase-indexer-engine	1.3-cdh5.0.3
	com.ngdata	hbase-indexer-model	1.3-cdh5.0.3
	com.ngdata	hbase-indexer-morphlines	1.3-cdh5.0.3
	com.ngdata	hbase-indexer-mr	1.3-cdh5.0.3
	com.ngdata	hbase-indexer-server	1.3-cdh5.0.3
	com.ngdata	hbase-sep-api	1.3-cdh5.0.3
	com.ngdata	hbase-sep-demo	1.3-cdh5.0.3
	com.ngdata	hbase-sep-impl-common	1.3-cdh5.0.3

Version and Download Information

Project	groupId	artifactId	version
	com.ngdata	hbase-sep-tools	1.3-cdh5.0.3

CDH 5.0.2

The following table lists the project name, groupId, artifactId, and version required to access each CDH 5 artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.0.2

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.0.2
Hadoop MRv1	org.apache.hadoop	hadoop-client	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-core	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-examples	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-minicluster	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-streaming	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-test	2.3.0-cdh5.0.2
	org.apache.hadoop	hadoop-tools	2.3.0-cdh5.0.2
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.0.2
	org.apache.hive	hive-beeline	0.12.0-cdh5.0.2
	org.apache.hive	hive-cli	0.12.0-cdh5.0.2
	org.apache.hive	hive-common	0.12.0-cdh5.0.2
	org.apache.hive	hive-contrib	0.12.0-cdh5.0.2
	org.apache.hive	hive-exec	0.12.0-cdh5.0.2
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.0.2
	org.apache.hive	hive-hwi	0.12.0-cdh5.0.2
	org.apache.hive	hive-jdbc	0.12.0-cdh5.0.2
	org.apache.hive	hive-metastore	0.12.0-cdh5.0.2
	org.apache.hive	hive-serde	0.12.0-cdh5.0.2
	org.apache.hive	hive-service	0.12.0-cdh5.0.2
	org.apache.hive	hive-shims	0.12.0-cdh5.0.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hive	hive-testutils	0.12.0-cdh5.0.2
	org.apache.hcatalog	hive-hcatalog-core	0.12.0-cdh5.0.2
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.0.2
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.0.2
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.0.2
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.0.2
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.0.2
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.0.2
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.0.2
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.0.2
Apache HBase	org.apache.hbase	hbase-client	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-common	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-examples	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-hadoop-compat	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-hadoop2-compat	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-it	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-prefix-tree	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-protocol	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-server	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-shell	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-testing-util	0.96.1.1-cdh5.0.2
	org.apache.hbase	hbase-thrift	0.96.1.1-cdh5.0.2
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.0.2
Apache Sqoop	org.apache.sqoop	sqoop	1.4.4-cdh5.0.2
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.0.2
	org.apache.pig	pigsmoke	0.12.0-cdh5.0.2
	org.apache.pig	pigunit	0.12.0-cdh5.0.2
Apache Flume	org.apache.flume	flume-ng-configuration	1.4.0-cdh5.0.2
	org.apache.flume	flume-ng-core	1.4.0-cdh5.0.2
	org.apache.flume	flume-ng-embedded-agent	1.4.0-cdh5.0.2

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-node	1.4.0-cdh5.0.2
	org.apache.flume	flume-ng-sdk	1.4.0-cdh5.0.2
	org.apache.flume	flume-ng-tests	1.4.0-cdh5.0.2
	org.apache.flume	flume-tools	1.4.0-cdh5.0.2
	org.apache.flume	flume-file-channel	1.4.0-cdh5.0.2
	org.apache.flume	flume-jdbc-channel	1.4.0-cdh5.0.2
	org.apache.flume	flume-ng-log4jappender	1.4.0-cdh5.0.2
	org.apache.flume	flume-avro-source	1.4.0-cdh5.0.2
	org.apache.flume	flume-thrift-source	1.4.0-cdh5.0.2
	org.apache.flume	flume-dataset-sink	1.4.0-cdh5.0.2
	org.apache.flume	flume-hdfs-sink	1.4.0-cdh5.0.2
	org.apache.flume	flume-irc-sink	1.4.0-cdh5.0.2
	org.apache.flume	flume-ng-elasticsearch-sink	1.4.0-cdh5.0.2
	org.apache.flume	flume-ng-hbase-sink	1.4.0-cdh5.0.2
	org.apache.flume	flume-ng-morphline-solr-sink	1.4.0-cdh5.0.2
	org.apache.flume	flume-jms-source	1.4.0-cdh5.0.2
	org.apache.flume	flume-scribe-source	1.4.0-cdh5.0.2
	org.apache.flume	flume-twitter-source	1.4.0-cdh5.0.2
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-core	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-examples	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.0.2

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.0.2
	org.apache.oozie	oozie-tools	4.0.0-cdh5.0.2
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.0.2
Apache Mahout	org.apache.mahout	mahout-buildtools	0.8-cdh5.0.2
	org.apache.mahout	mahout-core	0.8-cdh5.0.2
	org.apache.mahout	mahout-examples	0.8-cdh5.0.2
	org.apache.mahout	mahout-integration	0.8-cdh5.0.2
	org.apache.mahout	mahout-math	0.8-cdh5.0.2
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-chef	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-cli	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-core	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-druid	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-examples	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-hama	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-pig	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-solr	0.9.0-cdh5.0.2

Project	groupId	artifactId	version
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.0.2
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.0.2
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.0.2
Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.0.2
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.0.2
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-binding-solr	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-core-common	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-core-model-db	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-core-model-search	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-dist	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-policy-common	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-policy-db	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-policy-search	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-provider-common	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-provider-file	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-tests-hive	1.2.0-cdh5.0.2
	org.apache.sentry	sentry-tests-solr	1.2.0-cdh5.0.2
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.0.2
	com.twitter	parquet-cascading	1.2.5-cdh5.0.2
	com.twitter	parquet-column	1.2.5-cdh5.0.2

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-common	1.2.5-cdh5.0.2
	com.twitter	parquet-encoding	1.2.5-cdh5.0.2
	com.twitter	parquet-format	1.0.0-cdh5.0.2
	com.twitter	parquet-generator	1.2.5-cdh5.0.2
	com.twitter	parquet-hadoop	1.2.5-cdh5.0.2
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.0.2
	com.twitter	parquet-pig	1.2.5-cdh5.0.2
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.0.2
	com.twitter	parquet-scrooge	1.2.5-cdh5.0.2
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.0.2
	com.twitter	parquet-thrift	1.2.5-cdh5.0.2
Llama	com.doudera.llama	llama	1.0.0-cdh5.0.2
Apache Spark	org.apache.spark	spark-bagel_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-core_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-examples_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-graphx_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-mllib_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-repl_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-streaming-flume_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-streaming-kafka_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-streaming-mqtt_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-streaming-twitter_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-streaming-zeromq_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-streaming_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-tools_2.10	0.9.0-cdh5.0.2
	org.apache.spark	spark-yarn_2.10	0.9.0-cdh5.0.2
Apache Crunch	org.apache.crunch	crunch-archetype	0.9.0-cdh5.0.2
	org.apache.crunch	crunch-contrib	0.9.0-cdh5.0.2
	org.apache.crunch	crunch-core	0.9.0-cdh5.0.2
	org.apache.crunch	crunch-examples	0.9.0-cdh5.0.2

Project	groupId	artifactId	version
	org.apache.crunch	crunch-hbase	0.9.0-cdh5.0.2
	org.apache.crunch	crunch-scrunch	0.9.0-cdh5.0.2
	org.apache.crunch	crunch-spark	0.9.0-cdh5.0.2
	org.apache.crunch	crunch-test	0.9.0-cdh5.0.2
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.0.2
	org.apache.avro	avro-compiler	1.7.5-cdh5.0.2
	org.apache.avro	avro-ipc	1.7.5-cdh5.0.2
	org.apache.avro	avro-mapred	1.7.5-cdh5.0.2
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.0.2
	org.apache.avro	avro-protobuf	1.7.5-cdh5.0.2
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.0.2
	org.apache.avro	avro-thrift	1.7.5-cdh5.0.2
	org.apache.avro	avro-tools	1.7.5-cdh5.0.2
	org.apache.avro	trevni-avro	1.7.5-cdh5.0.2
	org.apache.avro	trevni-core	1.7.5-cdh5.0.2
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.0.2
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.0.2
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.0.2
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.0.2
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.0.2

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.0.2
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.0.2
	org.kitesdk	kite-tools	0.10.0-cdh5.0.2
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-classification	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-core	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-demo	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-facet	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-join	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-memory	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-misc	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-queries	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.0.2

Project	groupId	artifactId	version
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.0.2
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.0.2
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.0.2
	org.apache.solr	solr-cell	4.4.0-cdh5.0.2
	org.apache.solr	solr-clustering	4.4.0-cdh5.0.2
	org.apache.solr	solr-core	4.4.0-cdh5.0.2
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.0.2
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.0.2
	org.apache.solr	solr-langid	4.4.0-cdh5.0.2
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.0.2
	org.apache.solr	solr-solrj	4.4.0-cdh5.0.2
	org.apache.solr	solr-test-framework	4.4.0-cdh5.0.2
	org.apache.solr	solr-uima	4.4.0-cdh5.0.2
	org.apache.solr	solr-velocity	4.4.0-cdh5.0.2
Search	com.cloudera.search	search-mr	1.0.0-cdh5.0.2
HBase Indexer	com.ngdata	hbase-indexer-all	1.3-cdh5.0.2
	com.ngdata	hbase-indexer-cli	1.3-cdh5.0.2
	com.ngdata	hbase-indexer-common	1.3-cdh5.0.2
	com.ngdata	hbase-indexer-demo	1.3-cdh5.0.2
	com.ngdata	hbase-indexer-dist	1.3-cdh5.0.2
	com.ngdata	hbase-indexer-engine	1.3-cdh5.0.2
	com.ngdata	hbase-indexer-model	1.3-cdh5.0.2
	com.ngdata	hbase-indexer-morphlines	1.3-cdh5.0.2
	com.ngdata	hbase-indexer-mr	1.3-cdh5.0.2
	com.ngdata	hbase-indexer-server	1.3-cdh5.0.2
	com.ngdata	hbase-sep-api	1.3-cdh5.0.2
	com.ngdata	hbase-sep-demo	1.3-cdh5.0.2
	com.ngdata	hbase-sep-impl-common	1.3-cdh5.0.2
	com.ngdata	hbase-sep-tools	1.3-cdh5.0.2

Version and Download Information

CDH 5.0.1

The following table lists the project name, groupId, artifactId, and version required to access each CDH 5.0.1 artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.0.1

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.0.1
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.0.1
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	230-mr1-cdh5.0.1
	org.apache.hadoop	hadoop-core	230-mr1-cdh5.0.1
	org.apache.hadoop	hadoop-examples	230-mr1-cdh5.0.1
	org.apache.hadoop	hadoop-minicluster	230-mr1-cdh5.0.1
	org.apache.hadoop	hadoop-streaming	230-mr1-cdh5.0.1
	org.apache.hadoop	hadoop-test	230-mr1-cdh5.0.1
	org.apache.hadoop	hadoop-tools	230-mr1-cdh5.0.1
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.0.1
	org.apache.hive	hive-beeline	0.12.0-cdh5.0.1
	org.apache.hive	hive-cli	0.12.0-cdh5.0.1
	org.apache.hive	hive-common	0.12.0-cdh5.0.1
	org.apache.hive	hive-contrib	0.12.0-cdh5.0.1
	org.apache.hive	hive-exec	0.12.0-cdh5.0.1
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.0.1
	org.apache.hive	hive-hwi	0.12.0-cdh5.0.1
	org.apache.hive	hive-jdbc	0.12.0-cdh5.0.1
	org.apache.hive	hive-metastore	0.12.0-cdh5.0.1
	org.apache.hive	hive-serde	0.12.0-cdh5.0.1
	org.apache.hive	hive-service	0.12.0-cdh5.0.1
	org.apache.hive	hive-shims	0.12.0-cdh5.0.1
	org.apache.hive	hive-testutils	0.12.0-cdh5.0.1
	org.apache.hive	hive-hcatalog-core	0.12.0-cdh5.0.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.0.1
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.0.1
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.0.1
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.0.1
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.0.1
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.0.1
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.0.1
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.0.1
Apache HBase	org.apache.hbase	hbase-client	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-common	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-examples	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-hadoop-compat	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-hadoop2-compat	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-it	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-prefix-tree	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-protocol	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-server	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-shell	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-testing-util	0.96.1.1-cdh5.0.1
	org.apache.hbase	hbase-thrift	0.96.1.1-cdh5.0.1
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.0.1
Apache Sqoop 1	org.apache.sqoop	sqoop	1.4.4-cdh5.0.1
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.0.1
	org.apache.pig	pigsmoke	0.12.0-cdh5.0.1
	org.apache.pig	pigunit	0.12.0-cdh5.0.1
Apache Flume	org.apache.flume	flume-ng-configuration	1.4.0-cdh5.0.1
	org.apache.flume	flume-ng-core	1.4.0-cdh5.0.1
	org.apache.flume	flume-ng-embedded-agent	1.4.0-cdh5.0.1
	org.apache.flume	flume-ng-node	1.4.0-cdh5.0.1

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-sdk	1.4.0-cdh5.0.1
	org.apache.flume	flume-ng-tests	1.4.0-cdh5.0.1
	org.apache.flume	flume-tools	1.4.0-cdh5.0.1
	org.apache.flume	flume-file-channel	1.4.0-cdh5.0.1
	org.apache.flume	flume-jdbc-channel	1.4.0-cdh5.0.1
	org.apache.flume	flume-ng-log4jappender	1.4.0-cdh5.0.1
	org.apache.flume	flume-avro-source	1.4.0-cdh5.0.1
	org.apache.flume	flume-thrift-source	1.4.0-cdh5.0.1
	org.apache.flume	flume-dataset-sink	1.4.0-cdh5.0.1
	org.apache.flume	flume-hdfs-sink	1.4.0-cdh5.0.1
	org.apache.flume	flume-irc-sink	1.4.0-cdh5.0.1
	org.apache.flume	flume-ng-elasticsearch-sink	1.4.0-cdh5.0.1
	org.apache.flume	flume-ng-hbase-sink	1.4.0-cdh5.0.1
	org.apache.flume	flume-ng-morphline-solr-sink	1.4.0-cdh5.0.1
	org.apache.flume	flume-jms-source	1.4.0-cdh5.0.1
	org.apache.flume	flume-scribe-source	1.4.0-cdh5.0.1
	org.apache.flume	flume-twitter-source	1.4.0-cdh5.0.1
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-core	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-examples	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.0.1

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.0.1
	org.apache.oozie	oozie-tools	4.0.0-cdh5.0.1
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.0.1
Apache Mahout	org.apache.mahout	mahout-buildtools	0.8-cdh5.0.1
	org.apache.mahout	mahout-core	0.8-cdh5.0.1
	org.apache.mahout	mahout-examples	0.8-cdh5.0.1
	org.apache.mahout	mahout-integration	0.8-cdh5.0.1
	org.apache.mahout	mahout-math	0.8-cdh5.0.1
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-chef	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-cli	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-core	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-druid	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-examples	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-hama	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-pig	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-solr	0.9.0-cdh5.0.1
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.0.1

Project	groupId	artifactId	version
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.0.1
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.0.1
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.0.1
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.0.1
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-binding-solr	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-core-common	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-core-model-db	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-core-model-search	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-dist	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-policy-common	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-policy-db	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-policy-search	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-provider-common	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-provider-file	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-tests-hive	1.2.0-cdh5.0.1
	org.apache.sentry	sentry-tests-solr	1.2.0-cdh5.0.1
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.0.1
	com.twitter	parquet-cascading	1.2.5-cdh5.0.1
	com.twitter	parquet-column	1.2.5-cdh5.0.1
	com.twitter	parquet-common	1.2.5-cdh5.0.1

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-encoding	1.2.5-cdh5.0.1
	com.twitter	parquet-format	1.0.0-cdh5.0.1
	com.twitter	parquet-generator	1.2.5-cdh5.0.1
	com.twitter	parquet-hadoop	1.2.5-cdh5.0.1
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.0.1
	com.twitter	parquet-pig	1.2.5-cdh5.0.1
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.0.1
	com.twitter	parquet-scrooge	1.2.5-cdh5.0.1
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.0.1
	com.twitter	parquet-thrift	1.2.5-cdh5.0.1
Llama	com.doudera.llama	llama	1.0.0-cdh5.0.1
Apache Spark	org.apache.spark	spark-bagel_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-core_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-examples_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-graphx_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-mllib_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-repl_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-streaming-flume_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-streaming-kafka_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-streaming-mqtt_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-streaming-twitter_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-streaming-zeromq_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-streaming_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-tools_2.10	0.9.0-cdh5.0.1
	org.apache.spark	spark-yarn_2.10	0.9.0-cdh5.0.1
Apache Crunch	org.apache.crunch	crunch-archetype	0.9.0-cdh5.0.1
	org.apache.crunch	crunch-contrib	0.9.0-cdh5.0.1
	org.apache.crunch	crunch-core	0.9.0-cdh5.0.1
	org.apache.crunch	crunch-examples	0.9.0-cdh5.0.1
	org.apache.crunch	crunch-hbase	0.9.0-cdh5.0.1

Project	groupId	artifactId	version
	org.apache.crunch	crunch-scrunch	0.9.0-cdh5.0.1
	org.apache.crunch	crunch-spark	0.9.0-cdh5.0.1
	org.apache.crunch	crunch-test	0.9.0-cdh5.0.1
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.0.1
	org.apache.avro	avro-compiler	1.7.5-cdh5.0.1
	org.apache.avro	avro-ipc	1.7.5-cdh5.0.1
	org.apache.avro	avro-mapred	1.7.5-cdh5.0.1
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.0.1
	org.apache.avro	avro-protobuf	1.7.5-cdh5.0.1
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.0.1
	org.apache.avro	avro-thrift	1.7.5-cdh5.0.1
	org.apache.avro	avro-tools	1.7.5-cdh5.0.1
	org.apache.avro	trevni-avro	1.7.5-cdh5.0.1
	org.apache.avro	trevni-core	1.7.5-cdh5.0.1
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.0.1
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.0.1
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.0.1
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.0.1
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.0.1

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.0.1
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.0.1
	org.kitesdk	kite-tools	0.10.0-cdh5.0.1
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-classification	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-core	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-demo	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-facet	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-join	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-memory	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-misc	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-queries	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.0.1

Project	groupId	artifactId	version
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.0.1
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.0.1
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.0.1
	org.apache.solr	solr-cell	4.4.0-cdh5.0.1
	org.apache.solr	solr-clustering	4.4.0-cdh5.0.1
	org.apache.solr	solr-core	4.4.0-cdh5.0.1
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.0.1
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.0.1
	org.apache.solr	solr-langid	4.4.0-cdh5.0.1
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.0.1
	org.apache.solr	solr-solrj	4.4.0-cdh5.0.1
	org.apache.solr	solr-test-framework	4.4.0-cdh5.0.1
	org.apache.solr	solr-uima	4.4.0-cdh5.0.1
	org.apache.solr	solr-velocity	4.4.0-cdh5.0.1
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.0.1
Lily HBase Indexer	com.ngdata	hbase-indexer-all	1.3-cdh5.0.1
	com.ngdata	hbase-indexer-cli	1.3-cdh5.0.1
	com.ngdata	hbase-indexer-common	1.3-cdh5.0.1
	com.ngdata	hbase-indexer-demo	1.3-cdh5.0.1
	com.ngdata	hbase-indexer-dist	1.3-cdh5.0.1
	com.ngdata	hbase-indexer-engine	1.3-cdh5.0.1
	com.ngdata	hbase-indexer-model	1.3-cdh5.0.1
	com.ngdata	hbase-indexer-morphlines	1.3-cdh5.0.1
	com.ngdata	hbase-indexer-mr	1.3-cdh5.0.1
	com.ngdata	hbase-indexer-server	1.3-cdh5.0.1
	com.ngdata	hbase-sep-api	1.3-cdh5.0.1
	com.ngdata	hbase-sep-demo	1.3-cdh5.0.1
	com.ngdata	hbase-sep-impl-common	1.3-cdh5.0.1
	com.ngdata	hbase-sep-tools	1.3-cdh5.0.1

Version and Download Information

CDH 5.0.0

The following table lists the project name, groupId, artifactId, and version required to access each CDH 5.0.0 artifact.

Project	groupId	artifactId	version
Apache Hadoop	org.apache.hadoop	hadoop-annotations	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-archives	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-assemblies	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-auth	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-common	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-datajoin	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-distcp	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-extras	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-gridmix	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-hdfs	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-hdfs-nfs	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-mapreduce-client-app	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-mapreduce-client-common	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-mapreduce-client-core	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-mapreduce-client-hs	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-mapreduce-examples	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-maven-plugins	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-minikdc	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-nfs	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-openstack	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-rumen	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-sls	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-yarn-api	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	2.3.0-cdh5.0.0

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-client	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-yarn-common	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-yarn-server-common	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-yarn-server-nodemanager	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-yarn-server-tests	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-yarn-server-web-proxy	2.3.0-cdh5.0.0
	org.apache.hadoop	hadoop-hdfs-bkjournal	2.3.0-cdh5.0.0
Apache Hadoop MRv1	org.apache.hadoop	hadoop-client	230-mr1-cdh5.0.0
	org.apache.hadoop	hadoop-core	230-mr1-cdh5.0.0
	org.apache.hadoop	hadoop-examples	230-mr1-cdh5.0.0
	org.apache.hadoop	hadoop-minicluster	230-mr1-cdh5.0.0
	org.apache.hadoop	hadoop-streaming	230-mr1-cdh5.0.0
	org.apache.hadoop	hadoop-test	230-mr1-cdh5.0.0
	org.apache.hadoop	hadoop-tools	230-mr1-cdh5.0.0
Apache Hive	org.apache.hive	hive-ant	0.12.0-cdh5.0.0
	org.apache.hive	hive-beeline	0.12.0-cdh5.0.0
	org.apache.hive	hive-cli	0.12.0-cdh5.0.0
	org.apache.hive	hive-common	0.12.0-cdh5.0.0
	org.apache.hive	hive-contrib	0.12.0-cdh5.0.0
	org.apache.hive	hive-exec	0.12.0-cdh5.0.0
	org.apache.hive	hive-hbase-handler	0.12.0-cdh5.0.0
	org.apache.hive	hive-hwi	0.12.0-cdh5.0.0
	org.apache.hive	hive-jdbc	0.12.0-cdh5.0.0
	org.apache.hive	hive-metastore	0.12.0-cdh5.0.0
	org.apache.hive	hive-serde	0.12.0-cdh5.0.0
	org.apache.hive	hive-service	0.12.0-cdh5.0.0
	org.apache.hive	hive-shims	0.12.0-cdh5.0.0
	org.apache.hive	hive-testutils	0.12.0-cdh5.0.0
	org.apache.hive	hive-hcatalog-core	0.12.0-cdh5.0.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.hcatalog	hive-hcatalog-hbase-storage-handler	0.12.0-cdh5.0.0
	org.apache.hcatalog	hive-hcatalog-pig-adapter	0.12.0-cdh5.0.0
	org.apache.hcatalog	hive-hcatalog-server-extensions	0.12.0-cdh5.0.0
	org.apache.hcatalog	hive-webhcat	0.12.0-cdh5.0.0
	org.apache.hcatalog	hive-webhcat-java-client	0.12.0-cdh5.0.0
	org.apache.hiveshims	hive-shims-0.23	0.12.0-cdh5.0.0
	org.apache.hiveshims	hive-shims-common	0.12.0-cdh5.0.0
	org.apache.hiveshims	hive-shims-common-secure	0.12.0-cdh5.0.0
Apache HBase	org.apache.hbase	hbase-client	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-common	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-examples	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-hadoop-compat	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-hadoop2-compat	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-it	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-prefix-tree	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-protocol	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-server	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-shell	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-testing-util	0.96.1.1-cdh5.0.0
	org.apache.hbase	hbase-thrift	0.96.1.1-cdh5.0.0
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.4.5-cdh5.0.0
Apache Sqoop 1	org.apache.sqoop	sqoop	1.4.4-cdh5.0.0
Apache Pig	org.apache.pig	pig	0.12.0-cdh5.0.0
	org.apache.pig	pigsmoke	0.12.0-cdh5.0.0
	org.apache.pig	pigunit	0.12.0-cdh5.0.0
Apache Flume	org.apache.flume	flume-ng-configuration	1.4.0-cdh5.0.0
	org.apache.flume	flume-ng-core	1.4.0-cdh5.0.0
	org.apache.flume	flume-ng-embedded-agent	1.4.0-cdh5.0.0
	org.apache.flume	flume-ng-node	1.4.0-cdh5.0.0

Project	groupId	artifactId	version
	org.apache.flume	flume-ng-sdk	1.4.0-cdh5.0.0
	org.apache.flume	flume-ng-tests	1.4.0-cdh5.0.0
	org.apache.flume	flume-tools	1.4.0-cdh5.0.0
	org.apache.flume	flume-file-channel	1.4.0-cdh5.0.0
	org.apache.flume	flume-jdbc-channel	1.4.0-cdh5.0.0
	org.apache.flume	flume-ng-log4jappender	1.4.0-cdh5.0.0
	org.apache.flume	flume-avro-source	1.4.0-cdh5.0.0
	org.apache.flume	flume-thrift-source	1.4.0-cdh5.0.0
	org.apache.flume	flume-dataset-sink	1.4.0-cdh5.0.0
	org.apache.flume	flume-hdfs-sink	1.4.0-cdh5.0.0
	org.apache.flume	flume-irc-sink	1.4.0-cdh5.0.0
	org.apache.flume	flume-ng-elasticsearch-sink	1.4.0-cdh5.0.0
	org.apache.flume	flume-ng-hbase-sink	1.4.0-cdh5.0.0
	org.apache.flume	flume-ng-morphline-solr-sink	1.4.0-cdh5.0.0
	org.apache.flume	flume-jms-source	1.4.0-cdh5.0.0
	org.apache.flume	flume-scribe-source	1.4.0-cdh5.0.0
	org.apache.flume	flume-twitter-source	1.4.0-cdh5.0.0
Apache Oozie	org.apache.oozie	oozie-client	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-core	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-examples	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-hadoop	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-hadoop-distcp	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-hadoop-test	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-hadoop-utils	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-hbase	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-hcatalog	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-sharelib-distcp	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-sharelib-hcatalog	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-sharelib-hive	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-sharelib-hive2	4.0.0-cdh5.0.0

Version and Download Information

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-oozie	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-sharelib-pig	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-sharelib-sqoop	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-sharelib-streaming	4.0.0-cdh5.0.0
	org.apache.oozie	oozie-tools	4.0.0-cdh5.0.0
	org.apache.oozie.test	oozie-mini	4.0.0-cdh5.0.0
Apache Mahout	org.apache.mahout	mahout-buildtools	0.8-cdh5.0.0
	org.apache.mahout	mahout-core	0.8-cdh5.0.0
	org.apache.mahout	mahout-examples	0.8-cdh5.0.0
	org.apache.mahout	mahout-integration	0.8-cdh5.0.0
	org.apache.mahout	mahout-math	0.8-cdh5.0.0
Apache Whirr	org.apache.whirr	whirr-build-tools	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-cassandra	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-cdh	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-chef	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-cli	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-core	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-druid	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-elasticsearch	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-examples	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-ganglia	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-hadoop	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-hama	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-hbase	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-kerberos	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-mahout	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-pig	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-puppet	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-solr	0.9.0-cdh5.0.0
	org.apache.whirr	whirr-yarn	0.9.0-cdh5.0.0

Project	groupId	artifactId	version
	org.apache.whirr	whirr-zookeeper	0.9.0-cdh5.0.0
DataFu	com.linkedin.datafu	datafu	1.1.0-cdh5.0.0
Apache Sqoop2	org.apache.sqoop	connector-sdk	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-client	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-common	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-core	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-docs	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-shell	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-spi	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-tomcat	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-tools	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-connector-generic-jdbc	1.99.3-cdh5.0.0
	org.apache.sqoop	sqoop-repository-derby	1.99.3-cdh5.0.0
Apache Sentry (incubating)	org.apache.sentry	sentry-binding-hive	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-binding-solr	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-core-common	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-core-model-db	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-core-model-search	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-dist	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-policy-common	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-policy-db	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-policy-search	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-provider-common	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-provider-file	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-tests-hive	1.2.0-cdh5.0.0
	org.apache.sentry	sentry-tests-solr	1.2.0-cdh5.0.0
Parquet	com.twitter	parquet-avro	1.2.5-cdh5.0.0
	com.twitter	parquet-cascading	1.2.5-cdh5.0.0
	com.twitter	parquet-column	1.2.5-cdh5.0.0
	com.twitter	parquet-common	1.2.5-cdh5.0.0

Version and Download Information

Project	groupId	artifactId	version
	com.twitter	parquet-encoding	1.2.5-cdh5.0.0
	com.twitter	parquet-format	1.0.0-cdh5.0.0
	com.twitter	parquet-generator	1.2.5-cdh5.0.0
	com.twitter	parquet-hadoop	1.2.5-cdh5.0.0
	com.twitter	parquet-hadoop-bundle	1.2.5-cdh5.0.0
	com.twitter	parquet-pig	1.2.5-cdh5.0.0
	com.twitter	parquet-pig-bundle	1.2.5-cdh5.0.0
	com.twitter	parquet-scrooge	1.2.5-cdh5.0.0
	com.twitter	parquet-test-hadoop2	1.2.5-cdh5.0.0
	com.twitter	parquet-thrift	1.2.5-cdh5.0.0
Llama	com.doudera.llama	llama	1.0.0-cdh5.0.0
Apache Spark	org.apache.spark	spark-bagel_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-core_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-examples_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-graphx_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-mllib_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-repl_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-streaming-flume_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-streaming-kafka_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-streaming-mqtt_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-streaming-twitter_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-streaming-zeromq_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-streaming_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-tools_2.10	0.9.0-cdh5.0.0
	org.apache.spark	spark-yarn_2.10	0.9.0-cdh5.0.0
Apache Crunch	org.apache.crunch	crunch-archetype	0.9.0-cdh5.0.0
	org.apache.crunch	crunch-contrib	0.9.0-cdh5.0.0
	org.apache.crunch	crunch-core	0.9.0-cdh5.0.0
	org.apache.crunch	crunch-examples	0.9.0-cdh5.0.0
	org.apache.crunch	crunch-hbase	0.9.0-cdh5.0.0

Project	groupId	artifactId	version
	org.apache.crunch	crunch-scrunch	0.9.0-cdh5.0.0
	org.apache.crunch	crunch-spark	0.9.0-cdh5.0.0
	org.apache.crunch	crunch-test	0.9.0-cdh5.0.0
Apache Avro	org.apache.avro	avro	1.7.5-cdh5.0.0
	org.apache.avro	avro-compiler	1.7.5-cdh5.0.0
	org.apache.avro	avro-ipc	1.7.5-cdh5.0.0
	org.apache.avro	avro-mapred	1.7.5-cdh5.0.0
	org.apache.avro	avro-maven-plugin	1.7.5-cdh5.0.0
	org.apache.avro	avro-protobuf	1.7.5-cdh5.0.0
	org.apache.avro	avro-service-archetype	1.7.5-cdh5.0.0
	org.apache.avro	avro-thrift	1.7.5-cdh5.0.0
	org.apache.avro	avro-tools	1.7.5-cdh5.0.0
	org.apache.avro	trevni-avro	1.7.5-cdh5.0.0
	org.apache.avro	trevni-core	1.7.5-cdh5.0.0
Kite SDK	org.kitesdk	kite-data-core	0.10.0-cdh5.0.0
	org.kitesdk	kite-data-crunch	0.10.0-cdh5.0.0
	org.kitesdk	kite-data-hbase	0.10.0-cdh5.0.0
	org.kitesdk	kite-data-hcatalog	0.10.0-cdh5.0.0
	org.kitesdk	kite-maven-plugin	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-avro	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-core	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-hadoop-core	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-hadoop-parquet-avro	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-hadoop-rcfile	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-hadoop-sequencefile	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-json	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-maxmind	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-metrics-servlets	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-saxon	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-solr-cell	0.10.0-cdh5.0.0

Version and Download Information

Project	groupId	artifactId	version
	org.kitesdk	kite-morphlines-solr-core	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-tika-core	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-tika-decompress	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-twitter	0.10.0-cdh5.0.0
	org.kitesdk	kite-morphlines-useragent	0.10.0-cdh5.0.0
	org.kitesdk	kite-tools	0.10.0-cdh5.0.0
Apache Solr	org.apache.lucene	lucene-analyzers-common	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-analyzers-icu	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-analyzers-kuromoji	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-analyzers-morfologik	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-analyzers-phonetic	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-analyzers-smartcn	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-analyzers-stempel	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-analyzers-uima	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-benchmark	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-classification	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-codecs	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-core	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-demo	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-facet	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-grouping	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-highlighter	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-join	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-memory	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-misc	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-queries	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-queryparser	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-replicator	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-sandbox	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-spatial	4.4.0-cdh5.0.0

Project	groupId	artifactId	version
	org.apache.lucene	lucene-suggest	4.4.0-cdh5.0.0
	org.apache.lucene	lucene-test-framework	4.4.0-cdh5.0.0
	org.apache.solr	solr-analysis-extras	4.4.0-cdh5.0.0
	org.apache.solr	solr-cell	4.4.0-cdh5.0.0
	org.apache.solr	solr-clustering	4.4.0-cdh5.0.0
	org.apache.solr	solr-core	4.4.0-cdh5.0.0
	org.apache.solr	solr-dataimporthandler	4.4.0-cdh5.0.0
	org.apache.solr	solr-dataimporthandler-extras	4.4.0-cdh5.0.0
	org.apache.solr	solr-langid	4.4.0-cdh5.0.0
	org.apache.solr	solr-sentry-handlers	4.4.0-cdh5.0.0
	org.apache.solr	solr-solrj	4.4.0-cdh5.0.0
	org.apache.solr	solr-test-framework	4.4.0-cdh5.0.0
	org.apache.solr	solr-uima	4.4.0-cdh5.0.0
	org.apache.solr	solr-velocity	4.4.0-cdh5.0.0
Cloudera Search	com.cloudera.search	search-mr	1.0.0-cdh5.0.0
Lily HBase Indexer	com.ngdata	hbase-indexer-all	1.3-cdh5.0.0
	com.ngdata	hbase-indexer-cli	1.3-cdh5.0.0
	com.ngdata	hbase-indexer-common	1.3-cdh5.0.0
	com.ngdata	hbase-indexer-demo	1.3-cdh5.0.0
	com.ngdata	hbase-indexer-dist	1.3-cdh5.0.0
	com.ngdata	hbase-indexer-engine	1.3-cdh5.0.0
	com.ngdata	hbase-indexer-model	1.3-cdh5.0.0
	com.ngdata	hbase-indexer-morphlines	1.3-cdh5.0.0
	com.ngdata	hbase-indexer-mr	1.3-cdh5.0.0
	com.ngdata	hbase-indexer-server	1.3-cdh5.0.0
	com.ngdata	hbase-sep-api	1.3-cdh5.0.0
	com.ngdata	hbase-sep-demo	1.3-cdh5.0.0
	com.ngdata	hbase-sep-impl-common	1.3-cdh5.0.0
	com.ngdata	hbase-sep-tools	1.3-cdh5.0.0

Cloudera Navigator Key Trustee Server Version and Download Information

You can install Key Trustee Server using parcels in Cloudera Manager or using packages. Key Trustee Server 5.4 requires a Custom Service Descriptor (CSD) for parcel-based installations, but Key Trustee Server 5.5 and higher do not. See [Parcels](#) for more information on parcels, and [Custom Service Descriptor Files](#) for more information on CSDs.

See the following sections for information on Key Trustee Server packaging.

For repositories, replace `<username>` and `<password>` with your credentials. If you do not know your credentials, contact your Cloudera account representative.

Navigator Key Trustee Server 5.15

Navigator Key Trustee Server 5.15.0

RHEL, CentOS, and Oracle Linux

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-server-5.15.0
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.14

Navigator Key Trustee Server 5.14.0

RHEL, CentOS, and Oracle Linux

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-server-5.14.0
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.13

Navigator Key Trustee Server 5.13.0

RHEL, CentOS, and Oracle Linux

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-server-5.13.0
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.12

Navigator Key Trustee Server 5.12.0

RHEL, CentOS, and Oracle Linux

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-server-5.12.0
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.11

Navigator Key Trustee Server 5.11.0

RHEL, CentOS, and Oracle Linux

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-server-5.11.0
```

Version and Download Information

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.10

Navigator Key Trustee Server 5.10.1

RHEL, CentOS, and Oracle Linux

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-server-5.10.1
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.10.0

RHEL, CentOS, and Oracle Linux

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-server-5.10.0
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.9

Navigator Key Trustee Server 5.9.0

RHEL or CentOS

- Package repository tarball:

- Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.

- Package installation command:

```
yum install keytrustee-server-5.9.0
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.8

Navigator Key Trustee Server 5.8.0

RHEL or CentOS

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.

- Package installation command:

```
yum install keytrustee-server-5.8.0
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.7

Navigator Key Trustee Server 5.7.0

RHEL or CentOS

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.

- Package installation command:

```
yum install keytrustee-server-5.7.0
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.5

Navigator Key Trustee Server 5.5.2

RHEL or CentOS

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-server-5.5.2
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.5.0

RHEL or CentOS

- Package repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Packages** from the **Package or Parcel** drop-down menu, select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-server-5.5.0
```

- Parcel repository tarball:
 - Go to the Cloudera Navigator Key Trustee Server [download page](#). Select **Parcels** from the **Package or Parcel** drop-down menu and click **DOWNLOAD NOW**.
- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.4



Note: Cloudera Navigator Key Trustee Server 5.4 is the major release following Key Trustee Server 3.8, and is the first Key Trustee Server release to support parcel-based installation using Cloudera Manager. 5.4.3 and 5.4.9 are the only maintenance releases of Key Trustee Server 5.4.

Navigator Key Trustee Server 5.4.9

RHEL or CentOS

- Package repository:

```
https://<username>:<password>@archive.gazzang.com/redhat/stable/6/
```


- Package installation command:

```
yum install keytrustee-server-5.4.9
```

- Parcel repository:

```
https://<username>:<password>@archive.gazzang.com/parcels/cloudera/keytrustee-server/5.4.9/stable/
```

- Custom Service Descriptor:

```
https://<username>:<password>@archive.gazzang.com/parcels/cloudera/keytrustee-server/5.4.9/stable/csd/
```

- Release notes:

[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.4.3

RHEL or CentOS

- Package repository:

```
https://<username>:<password>@archive.gazzang.com/redhat/stable/6/
```

- Package installation command:

```
yum install keytrustee-server-5.4.3
```

- Parcel repository:

```
https://<username>:<password>@archive.gazzang.com/parcels/cloudera/keytrustee-server/5.4.3/stable/
```

- Custom Service Descriptor:

```
https://<username>:<password>@archive.gazzang.com/parcels/cloudera/keytrustee-server/5.4.3/stable/csd/
```

- Release notes:

[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Navigator Key Trustee Server 5.4.0

RHEL or CentOS

- Package repository:

```
https://<username>:<password>@archive.gazzang.com/redhat/stable/6/
```

- Package installation command:

```
yum install keytrustee-server-5.4.0
```

- Parcel repository:

```
https://<username>:<password>@archive.gazzang.com/parcels/cloudera/keytrustee-server/5.4.0/stable/latest/
```

- Custom Service Descriptor:

```
https://<username>:<password>@archive.gazzang.com/parcels/cloudera/keytrustee-server/5.4.0/stable/latest/csd/
```

Version and Download Information

- Release notes:
[Cloudera Navigator Key Trustee Server Release Notes](#) on page 2785

Cloudera Navigator Key HSM Version and Download Information

Cloudera Navigator Key HSM 1.11

Cloudera Navigator Key HSM 1.11.0

RHEL or CentOS

- Package repository tarball:
 - Visit the Cloudera Navigator Key HSM [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-keyhsm-1.11.0
```

- Release notes:
[Cloudera Navigator Key HSM Release Notes](#) on page 2797

Cloudera Navigator Key HSM 1.10

Cloudera Navigator Key HSM 1.10.0

RHEL or CentOS

- Package repository tarball:
 - Visit the Cloudera Navigator Key HSM [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-keyhsm-1.10.0
```

- Release notes:
[Cloudera Navigator Key HSM Release Notes](#) on page 2797

Cloudera Navigator Key HSM 1.9

Cloudera Navigator Key HSM 1.9.0

RHEL or CentOS

- Package repository tarball:
 - Visit the Cloudera Navigator Key HSM [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-keyhsm-1.9.0
```

- Release notes:
[Cloudera Navigator Key HSM Release Notes](#) on page 2797

Cloudera Navigator Key HSM 1.8

Cloudera Navigator Key HSM 1.8.0

RHEL or CentOS

- Package repository tarball:
 - Visit the Cloudera Navigator Key HSM [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-keyhsm-1.8.0
```

- Release notes:
[Cloudera Navigator Key HSM Release Notes](#) on page 2797

Cloudera Navigator Key HSM 1.7

Cloudera Navigator Key HSM 1.7.0

RHEL or CentOS

- Package repository tarball:
 - Visit the Cloudera Navigator Key HSM [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-keyhsm-1.7.0
```

- Release notes:
[Cloudera Navigator Key HSM Release Notes](#) on page 2797

Cloudera Navigator Key HSM 1.6

Cloudera Navigator Key HSM 1.6.0

RHEL or CentOS

- Package repository tarball:
 - Visit the Cloudera Navigator Key HSM [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install keytrustee-keyhsm-1.6.0
```

- Release notes:
[Cloudera Navigator Key HSM Release Notes](#) on page 2797

Cloudera Navigator Key HSM 1.5

Cloudera Navigator Key HSM 1.5.1

RHEL or CentOS

- Package repository tarball:

Version and Download Information

- Visit the Cloudera Navigator Key HSM [download page](#). Select your OS version, and click **DOWNLOAD NOW**.

- Package installation command:

```
yum install keytrustee-keyhsm-1.5.1
```

- Release notes:

[Cloudera Navigator Key HSM Release Notes](#) on page 2797

Cloudera Navigator Key HSM 1.5.0

RHEL or CentOS

- Package repository tarball:

- Visit the Cloudera Navigator Key HSM [download page](#). Select your OS version, and click **DOWNLOAD NOW**.

- Package installation command:

```
yum install keytrustee-keyhsm-1.5.0
```

- Release notes:

[Cloudera Navigator Key HSM Release Notes](#) on page 2797

Cloudera Navigator Key HSM 1.4



Note: There are no Key HSM 1.4 maintenance releases.

Cloudera Navigator Key HSM 1.4.0

RHEL or CentOS

- Package repository:

```
https://<username>:<password>@archive.gazzang.com/redhat/stable/6/
```

- Package installation command:

```
yum install keytrustee-keyhsm-1.4.0
```

- Release notes:

[Cloudera Navigator Key HSM Release Notes](#) on page 2797

Key Trustee KMS Version and Download Information

Key Trustee KMS 5.16

Key Trustee KMS 5.16.0

- Parcel repository tarball:

- Visit the Key Trustee KMS [download page](#).
- Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.

- Click **DOWNLOAD NOW**.
- Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
[Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.15

Key Trustee KMS 5.15.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
[Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.14

Key Trustee KMS 5.14.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.

Version and Download Information

- Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.13

Key Trustee KMS 5.13.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.12.1

Key Trustee KMS 5.12.1

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).

- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.12

Key Trustee KMS 5.12.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.11

Key Trustee KMS 5.11.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:

Version and Download Information

- Visit the Key Trustee KMS [download page](#).
- Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
- Select your operating system from the **SELECT AN OS** drop-down menu.
- Click **DOWNLOAD NOW**.
- Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
[Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.10

Key Trustee KMS 5.10.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
[Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.9

Key Trustee KMS 5.9.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.

- Select your operating system from the **SELECT AN OS** drop-down menu.
- Click **DOWNLOAD NOW**.
- Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
[Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.8

Key Trustee KMS 5.8.2

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
[Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.8.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).

Version and Download Information

- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.7

Key Trustee KMS 5.7.4

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.7.1

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:

[Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.7.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:

[Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.5

Key Trustee KMS 5.5.4

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:

[Key Trustee KMS Release Notes](#) on page 2801.

Version and Download Information

Key Trustee KMS 5.5.0

- Parcel repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the Key Trustee KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801.

Key Trustee KMS 5.4

Key Trustee KMS 5.4.3

RHEL compatible

- Parcel repository:
 - <https://archive.cloudera.com/navigator-keytrustee5/parcels/5.4.3/>
- Package repository .repo file:

```
https://archive.cloudera.com/navigator-keytrustee5/redhat/<version>/x86_64/navigator-keytrustee/navigator-keytrustee5.repo
```

Replace *<version>* with 5 or 6, depending on your OS version.

Edit the `navigator-keytrustee5.repo` file and change the 5 at the end of the line that begins `baseurl=` to 5.4.3. For example:

```
baseurl=https://archive.cloudera.com/navigator-keytrustee5/redhat/<version>/x86_64/navigator-keytrustee/5.4.3/
```

- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801

SLES

- Parcel repository:
 - <https://archive.cloudera.com/navigator-keytrustee5/parcels/5.4.3/>

- Package repository:

```
https://archive.cloudera.com/navigator-keytrustee5/sles/11/x86_64/navigator-keytrustee/navigator-keytrustee5.repo
```

Edit the `navigator-keytrustee5.repo` file and change the 5 at the end of the line that begins `baseurl=` to 5.4.3. For example:

```
baseurl=https://archive.cloudera.com/navigator-keytrustee5/sles/11/x86_64/navigator-keytrustee/5.4.3/
```

- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801

Ubuntu

- Parcel repository:
 - <https://archive.cloudera.com/navigator-keytrustee5/parcels/5.4.3/>
- Package repository:

```
https://archive.cloudera.com/navigator-keytrustee5/ubuntu/<version>/amd64/navigator-keytrustee/cloudera.list
```

Replace `<version>` with `precise` for Ubuntu 12.04 or `trusty` for Ubuntu 14.04.

Edit the `cloudera.list` file and change `<distribution>-navigator-keytrustee5` to `<distribution>-navigator-keytrustee5.4.3` on both lines. For example, for Ubuntu Precise:

```
deb [arch=amd64]
https://archive.cloudera.com/navigator-keytrustee5/ubuntu/precise/amd64/navigator-keytrustee
precise-navigator-keytrustee5.4.3 contrib
deb-src
https://archive.cloudera.com/navigator-keytrustee5/ubuntu/precise/amd64/navigator-keytrustee
precise-navigator-keytrustee5.4.3 contrib
```

- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
 - [Key Trustee KMS Release Notes](#) on page 2801

Debian

- Parcel repository:
 - <https://archive.cloudera.com/navigator-keytrustee5/parcels/5.4.3/>
- Package repository:

```
https://archive.cloudera.com/navigator-keytrustee5/debian/wheezy/amd64/navigator-keytrustee/cloudera.list
```

Edit the `cloudera.list` file and change `wheezy-navigator-keytrustee5` to `wheezy-navigator-keytrustee5.4.3` on both lines. For example:

```
deb
https://archive.cloudera.com/navigator-keytrustee5/debian/wheezy/amd64/navigator-keytrustee
wheezy-navigator-keytrustee5.4.3 contrib
deb-src
```

Version and Download Information

```
https://archive.cloudera.com/navigator-keytrustee5/debian/wheezy/amd64/navigator-keytrustee-wheezy-navigator-keytrustee5.4.3 contrib
```

- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
[Key Trustee KMS Release Notes](#) on page 2801

Key Trustee KMS 5.4.0

RHEL compatible

- Parcel repository:
 - <https://archive.cloudera.com/navigator-keytrustee5/parcels/5.4.0/>
- Package repository .repo file:

```
https://archive.cloudera.com/navigator-keytrustee5/redhat/<version>/x86_64/navigator-keytrustee/navigator-keytrustee5.repo
```

Replace `<version>` with 5 or 6, depending on your OS version.

Edit the `navigator-keytrustee5.repo` file and change the 5 at the end of the line that begins `baseurl=` to 5.4.0. For example:

```
baseurl=https://archive.cloudera.com/navigator-keytrustee5/redhat/<version>/x86_64/navigator-keytrustee/5.4.0/
```

- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
[Key Trustee KMS Release Notes](#) on page 2801

SLES

- Parcel repository:
 - <https://archive.cloudera.com/navigator-keytrustee5/parcels/5.4.0/>
- Package repository:

```
https://archive.cloudera.com/navigator-keytrustee5/sles/11/x86_64/navigator-keytrustee/navigator-keytrustee5.repo
```

Edit the `navigator-keytrustee5.repo` file and change the 5 at the end of the line that begins `baseurl=` to 5.4.0. For example:

```
baseurl=https://archive.cloudera.com/navigator-keytrustee5/sles/11/x86_64/navigator-keytrustee/5.4.0/
```

- Installation instructions:
 - [Installing Key Trustee KMS](#)
- Release notes:
[Key Trustee KMS Release Notes](#) on page 2801

Ubuntu

- Parcel repository:

- <https://archive.cloudera.com/navigator-keytrustee5/parcels/5.4.0/>

- Package repository:

```
https://archive.cloudera.com/navigator-keytrustee5/ubuntu/<version>/amd64/navigator-keytrustee/cloudera.list
```

Replace `<version>` with `precise` for Ubuntu 12.04 or `trusty` for Ubuntu 14.04.

Edit the `cloudera.list` file and change `<distribution>-navigator-keytrustee5` to `<distribution>-navigator-keytrustee5.4.0` on both lines. For example, for Ubuntu Precise:

```
deb [arch=amd64]
https://archive.cloudera.com/navigator-keytrustee5/ubuntu/precise/amd64/navigator-keytrustee
precise-navigator-keytrustee5.4.0 contrib
deb-src
https://archive.cloudera.com/navigator-keytrustee5/ubuntu/precise/amd64/navigator-keytrustee
precise-navigator-keytrustee5.4.0 contrib
```

- Installation instructions:

- [Installing Key Trustee KMS](#)

- Release notes:

[Key Trustee KMS Release Notes](#) on page 2801

Debian

- Parcel repository:

- <https://archive.cloudera.com/navigator-keytrustee5/parcels/5.4.0/>

- Package repository:

```
https://archive.cloudera.com/navigator-keytrustee5/debian/wheezy/amd64/navigator-keytrustee/cloudera.list
```

Edit the `cloudera.list` file and change `wheezy-navigator-keytrustee5` to `wheezy-navigator-keytrustee5.4.0` on both lines. For example:

```
deb
https://archive.cloudera.com/navigator-keytrustee5/debian/wheezy/amd64/navigator-keytrustee
wheezy-navigator-keytrustee5.4.0 contrib
deb-src
https://archive.cloudera.com/navigator-keytrustee5/debian/wheezy/amd64/navigator-keytrustee
wheezy-navigator-keytrustee5.4.0 contrib
```

- Installation instructions:

- [Installing Key Trustee KMS](#)

- Release notes:

[Key Trustee KMS Release Notes](#) on page 2801

HSM KMS Version and Download Information

HSM KMS 5.15

HSM KMS 5.15.0

- Parcel repository tarball:

- Visit the HSM KMS [download page](#).

Version and Download Information

- Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
- Click **DOWNLOAD NOW**.
- Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the HSM KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Navigator HSM KMS Backed by Thales HSM](#)
 - [Installing Navigator HSM KMS Backed by Luna HSM](#)
- Release notes:
[Cloudera Navigator HSM KMS Release Notes](#) on page 2809

HSM KMS 5.14

HSM KMS 5.14.0

- Parcel repository tarball:
 - Visit the HSM KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the HSM KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Navigator HSM KMS Backed by Thales HSM](#)
 - [Installing Navigator HSM KMS Backed by Luna HSM](#)
- Release notes:
[Cloudera Navigator HSM KMS Release Notes](#) on page 2809

HSM KMS 5.13

HSM KMS 5.13.0

- Parcel repository tarball:
 - Visit the HSM KMS [download page](#).

- Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
- Click **DOWNLOAD NOW**.
- Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the HSM KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Navigator HSM KMS Backed by Thales HSM](#)
 - [Installing Navigator HSM KMS Backed by Luna HSM](#)
- Release notes:
 - [Cloudera Navigator HSM KMS Release Notes](#) on page 2809

HSM KMS 5.12.1

HSM KMS 5.12.1

- Parcel repository tarball:
 - Visit the HSM KMS [download page](#).
 - Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the HSM KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Navigator HSM KMS Backed by Thales HSM](#)
 - [Installing Navigator HSM KMS Backed by Luna HSM](#)
- Release notes:
 - [Cloudera Navigator HSM KMS Release Notes](#) on page 2809

HSM KMS 5.12

HSM KMS 5.12.0

- Parcel repository tarball:
 - Visit the HSM KMS [download page](#).

Version and Download Information

- Select **Parcels** from the **SELECT DOWNLOAD TYPE** drop-down menu.
- Click **DOWNLOAD NOW**.
- Set up an internal repository, following the instructions in [Creating and Using a Remote Parcel Repository for Cloudera Manager](#).
- Package repository tarball:
 - Visit the HSM KMS [download page](#).
 - Select **Packages** from the **SELECT DOWNLOAD TYPE** drop-down menu.
 - Select your operating system from the **SELECT AN OS** drop-down menu.
 - Click **DOWNLOAD NOW**.
 - Set up an internal repository, following the instructions in [Creating and Using a Package Repository for Cloudera Manager](#).
- Installation instructions:
 - [Installing Navigator HSM KMS Backed by Thales HSM](#)
 - [Installing Navigator HSM KMS Backed by Luna HSM](#)
- Release notes:
[Cloudera Navigator HSM KMS Release Notes](#) on page 2809

Cloudera Navigator Encrypt Version and Download Information

For repositories, replace `<username>` and `<password>` with your credentials. If you do not know your credentials, contact your Cloudera account team.

Navigator Encrypt 3.16

Navigator Encrypt 3.16.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Visit the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install navencrypt-3.16.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
zypper install navencrypt-3.16.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.

- Package installation command:

```
apt-get install navencrypt-3.16.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.16.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.15

Navigator Encrypt 3.15.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Visit the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install navencrypt-3.15.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
zypper install navencrypt-3.15.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.15.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.

Version and Download Information

- Package installation command:

```
apt-get install navencrypt-3.15.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.14

Navigator Encrypt 3.14.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Visit the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install navencrypt-3.14.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
zypper install navencrypt-3.14.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.14.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.14.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.13

Navigator Encrypt 3.13.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Visit the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install navencrypt-3.13.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
zypper install navencrypt-3.13.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.13.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.13.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.12

Navigator Encrypt 3.12.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Visit the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.

Version and Download Information

- Package installation command:

```
yum install navencrypt-3.12.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
zypper install navencrypt-3.12.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.12.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.12.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.11

Navigator Encrypt 3.11.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Visit the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install navencrypt-3.11.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.

- Package installation command:

```
zypper install navencrypt-3.11.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.11.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.11.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.10

Navigator Encrypt 3.10.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Visit the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install navencrypt-3.10.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
zypper install navencrypt-3.10.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.

Version and Download Information

- Package installation command:

```
apt-get install navencrypt-3.10.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.10.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.9

Navigator Encrypt 3.9.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Visit the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install navencrypt-3.9.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
zypper install navencrypt-3.9.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.9.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.

- Package installation command:

```
apt-get install navencrypt-3.9.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.8

Navigator Encrypt 3.8.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Visit the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install navencrypt-3.8.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
zypper install navencrypt-3.8.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.8.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.8.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.7

Navigator Encrypt 3.7.1

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository tarball:
 - Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
yum install navencrypt-3.7.1
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
zypper install navencrypt-3.7.1
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.7.1
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Go to the Cloudera Navigator Encrypt [download page](#). Select your OS version, and click **DOWNLOAD NOW**.
- Package installation command:

```
apt-get install navencrypt-3.7.1
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Navigator Encrypt 3.7.0

RHEL, CentOS, and Oracle Enterprise Linux

- Package repository:

```
https://<username>:<password>@archive.gazzang.com/redhat/stable/6/
```

- Package installation command:

```
yum install navencrypt-3.7.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

SLES

- Package repository:

```
https://<username>:<password>@archive.gazzang.com/sles/stable/11/
```

- Package installation command:

```
zypper install navencrypt-3.7.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Ubuntu

- Package repository:

```
https://<username>:<password>@archive.gazzang.com/ubuntu/stable/
```

- Package installation command:

```
apt-get install navencrypt-3.7.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Debian

- Package repository:

```
https://<username>:<password>@archive.gazzang.com/debian/stable/
```

- Package installation command:

```
apt-get install navencrypt-3.7.0
```

- Release notes:
[Cloudera Navigator Encrypt Release Notes](#) on page 2813

Impala Version and Download Information

The information in this section applies to CDH 4 clusters, where Impala is downloaded and installed separately from CDH itself. You can disregard this section if you are running a CDH 5 cluster. This information is included with the CDH 5 documentation for users who manage CDH 4 clusters through Cloudera Manager 5.

See the following sections for Impala download information:

- [Impala 2.2.0 and Higher](#) on page 1792 (latest 2.2.x, for CDH 5 only)
- [Impala 2.1.0](#) on page 1792 (latest 2.1.x, for CDH 4 or CDH 5)

Version and Download Information

- [Impala 2.0.1](#) on page 1793 (latest 2.0.x, for CDH 4 or CDH 5)
- [Impala 2.0.0](#) on page 1793
- [Impala 1.4.3](#) on page 1794 (latest 1.4.x, for CDH 4 or CDH 5)
- [Impala 1.4.0](#) on page 1794
- [Impala 1.3.1](#) on page 1795 (latest 1.3.x, for CDH 4 or CDH 5)
- [Impala 1.2.4](#) on page 1795 (latest 1.2.x, for CDH 4)
- [Impala 1.2.3](#) on page 1796 (for CDH 4; also bundled with CDH 5 beta 2)
- [Impala 1.2.2](#) on page 1797
- [Impala 1.2.1](#) on page 1797
- [Impala 1.2.0 \(Beta\)](#) on page 1798 (for use with CDH 5 beta 1)
- [Impala 1.1.1](#) on page 1798 (latest 1.1.x)
- [Impala 1.0.1](#) on page 1798 (latest 1.0.x)

For installation instructions, see [Setting Up Apache Impala Using the Command Line](#).

For information about new features in the latest release, and issues that are fixed or still outstanding, see the [Cloudera Release Guide](#).

Impala Download Information



Note: All packages are 64-bit.

Impala 2.2.0 and Higher



Important: Impala 2.2.0 and higher are only available for CDH 5, not for CDH 4. For version and compatibility information for the Impala bundled with CDH 5, see [CDH Version and Packaging Information](#).

Impala 2.1.0

Release Date: December 2014 **Status:** Production

These Impala packages work with CDH 4. For the Impala bundled with CDH 5, see [CDH Version and Packaging Information](#).

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/2/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/2/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/2/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list

Repository Type	Location	Repo or List file
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 2.0.1

Release Date: December 2014 **Status:** Production

These Impala packages work with CDH 4. For the Impala bundled with CDH 5, see [CDH Version and Packaging Information](#).

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/2.0.1/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/2.0.1/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/2.0.1/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 2.0.0

Release Date: October 2014 **Status:** Production

These Impala packages work with CDH 4. For the Impala bundled with CDH 5, see [CDH Version and Packaging Information](#).

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/2/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/2/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/2/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo

Version and Download Information

Repository Type	Location	Repo or List file
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 1.4.3

Release Date: December 2014 **Status:** Production

These Impala packages work with CDH 4. For the Impala bundled with CDH 5, see [CDH Version and Packaging Information](#).

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/1.4.3/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/1.4.3/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/1.4.3/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 1.4.0

Release Date: July 2014 **Status:** Production

These Impala packages work with CDH 4. For the Impala bundled with CDH 5, see [CDH Version and Packaging Information](#).

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/1.4.0/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo

Repository Type	Location	Repo or List file
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/1.4.0/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/1.4.0/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 1.3.1

Release Date: May 2014 **Status:** Production

These Impala packages work with CDH 4. For the Impala bundled with CDH 5, see [CDH Version and Packaging Information](#).

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/1.3.1/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/1.3.1/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/1.3.1/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 1.2.4

Release Date: February 2014 **Status:** Production

This Impala version works with CDH 4.

Version and Download Information

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/1.2.4/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/1.2.4/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/1.2.4/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 1.2.3

Release Date: December 2013 (for CDH 4), February 2013 (for CDH 5 beta 2) **Status:** Production

This Impala version works with CDH 4, and also comes bundled with CDH 5 beta 2.

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/1.2.3/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/1.2.3/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/1.2.3/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 1.2.2

Release Date: December 2013 **Status:** Superseded by Impala 1.2.3 due to important bug fix

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/1.2.2/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/1.2.2/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/1.2.2/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 1.2.1

Release Date: November 2013 **Status:** Production

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/1.2.1/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/1.2.1/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/1.2.1/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Version and Download Information

Impala 1.2.0 (Beta)

Release Date: October 2013 **Status:** Beta, for use with CDH 5 Beta 1

CDH 5 Project	Package Version	Tarball Version	Release Notes	Changes File
Impala	impala-1.2.0+0	(none)	here	here

Impala 1.1.1

Release Date: July 2013 **Status:** Production

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/1.1.1/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/1.1.1/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/1.1.1/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

Impala 1.0.1

Release Date: June 2013 **Status:** Production

Repository Type	Location	Repo or List file
Yum RHEL 6/CentOS 6 (64-bit)	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/1.0.1/	https://archive.cloudera.com/impala/redhat/6/x86_64/impala/cloudera-impala.repo
Yum RHEL 5/CentOS 5 (64-bit)	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/1.0.1/	https://archive.cloudera.com/impala/redhat/5/x86_64/impala/cloudera-impala.repo
Zypper SLES 11 (64-bit)	https://archive.cloudera.com/impala/sles/11/x86_64/impala/1.0.1/	https://archive.cloudera.com/impala/sles/11/x86_64/impala/cloudera-impala.repo
Apt-Get Ubuntu 10.04 (Lucid) (64-bit)	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/lucid/amd64/impala/cloudera.list

Repository Type	Location	Repo or List file
Apt-Get Ubuntu 12.04 (Precise) (64-bit)	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/	https://archive.cloudera.com/impala/ubuntu/precise/amd64/impala/cloudera.list
Apt-Get Debian (Squeeze) (64-bit)	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/	https://archive.cloudera.com/impala/debian/squeeze/amd64/impala/cloudera.list

What's New in Cloudera Documentation

This page describes new topics added and changes made to Cloudera documentation.

What's New in Cloudera Documentation in May, 2019

This section describes new topics added and major changes made to Cloudera documentation in May, 2019:

Table 43:

Product	What's New	Link
Cloudera Navigator	Improved information on <ul style="list-style-type: none"> Navigator TLS configuration Configuring Cloudera Navigator for LDAP 	<ul style="list-style-type: none"> Configuring TLS/SSL for Navigator Metadata Server Configuring Cloudera Navigator for LDAP

What's New in Cloudera Documentation in April, 2019

This section describes new topics added and major changes made to Cloudera documentation in April, 2019:

Table 44:

Product	What's New	Link
Cloudera Navigator	Added detailed information on how metadata is extracted from multiple sources and combined to generate lineage for data assets, specifically for Hive. This section helps understand when you can expect lineage to appear after data assets are created on the cluster.	Metadata Extraction Timing

What's New in Cloudera Documentation in March, 2019

This section describes new topics added and major changes made to Cloudera documentation in March, 2019:

Table 45:

Product	What's New	Link
Cloudera Data Science Workbench	Published a new Security Overview for Cloudera Data Science Workbench. This topic goes over the basics of the CDSW security model, the wildcard DNS requirement, and how authentication, authorization, and wire encryption work in Cloudera Data Science Workbench.	CDSW Security Overview

What's New in Cloudera Documentation in February, 2019

This section describes new topics added and major changes made to Cloudera documentation in February, 2019:

Table 46:

Product	What's New	Link
Apache Impala	Added a new section on load-balancing proxy in TLS-enabled cluster.	Special Proxy Considerations for TLS/SSL-Enabled Clusters
	Added a new section on enabling LDAP in Cloudera Manager.	Enabling LDAP in Cloudera Manager
	Updated docs to reflect decoupling of compute and storage in Hadoop clusters.	Components of the Impala Server
Apache Kudu	Added a recommendation to use nscd (name service caching daemon) for all name resolutions.	Slow Name Resolution and nscd
Hue	The Hue Guide has been completely re-organized and additional content added. Look for further improvements in the areas of performance tuning and reference architectures to be released soon.	<ul style="list-style-type: none"> • CDH 5.16 Hue Guide • CDH 5.15 Hue Guide

What's New in Cloudera Documentation in January, 2019

This section describes new topics added and major changes made to Cloudera documentation in January, 2019:

Table 47:

Product	What's New	Link
Apache Spark ML	Learn how to configure Spark ML to use native math libraries that accelerate model training speed for algorithms like Alternating Least Squares (ALS).	Using Native Math Libraries to Accelerate Spark Machine Learning Applications
Workload XM	The new Workload View feature enables you to break down workloads by specific criteria to perform deep-dive analysis on the queries. For example, you can use the Workload View feature to determine which users are executing workloads that do not adhere to SLAs. You can also examine how queries being sent to specific databases or that use specific pools are performing against SLAs.	<ul style="list-style-type: none"> • Classifying Workloads for Analysis with Workload Views • Video: Classifying Workloads to Gain Insights

What's New in Cloudera Documentation in November, 2018

This section describes new topics added and major changes made to Cloudera documentation in November, 2018:

Table 48:

Product	What's New	Link
Apache HBase	Information about how to move the HBase Master role from one host to another.	Moving HBase Master Role to Another Host
Cloudera Navigator Key Trustee KMS	Added a new procedure that describes how to move a Key Trustee KMS proxy service role instance from an existing cluster host to another cluster host. This feature is for 5.16.1 and later <i>only</i> .	Migrating a Key Trustee KMS Server Role Instance to a New Host
Workload Experience Manager (Workload XM)	With the release of Cloudera Manager 5.16.1, new functionality has been added to Workload XM to redact logs and queries and for proxy server support. These new features can be enabled by configuring the Telemetry Publisher service in Cloudera Manager. In addition to these new features, now you can download the SQL commands to address "Corrupt Table Statistics" and "Missing Table Statistics" query health checks and numerous usability enhancements have also been added.	What's New from Workload XM
	Information about how to configure Workload XM to tunnel through a firewall in your environment.	Configuring a Firewall for Workload XM
	Detailed description of the diagnostic data collection performed by Workload XM.	Workload XM Diagnostic Data Collection

What's New in Cloudera Documentation in October, 2018

This section describes new topics added and major changes made to Cloudera documentation in October, 2018:

Table 49:

Product	What's New	Link
Cloudera Data Science Workbench	Released Cloudera Data Science Workbench 1.4.2. This release fixes some critical bugs, including TSB-346: Risk of Data Loss on Cloudera Data Science Workbench Shutdown and Restart . Please read the TSB and Upgrade Notes carefully before you start upgrading or perform a shutdown/restart operation on any previous version of CDSW.	<ul style="list-style-type: none"> • Release Notes • Known Issues and Limitations
Apache HDFS	Documented how to enable authorization for HDFS web UIs.	Enabling Authorization for HDFS Web UIs
Apache Impala	Documented using a query option to set an execution time limit on queries.	Setting Time Limits on Long Running Queries

Product	What's New	Link
	Documented the scheduler-related query hints and options support Kudu tablets.	<ul style="list-style-type: none"> • Impala Hints • Impala Query Option
Apache Kudu	Noted that it is better to let Kudu manage its own striping over multiple devices rather than delegating the striping to a RAID-0 array.	Kudu Configuration
Cloudera Navigator	Added detailed steps for streaming Navigator audit events to a Kafka topic.	Publishing Audit Events to Kafka
	Added tips on how to get the most from your Navigator Audit Server implementation, including some maintenance steps that will help you make sure you are collecting the right audit events.	Maintaining Navigator Audit Server
Apache Sentry	Updated the Sentry privilege tables for Hive and Impala. They now include all possible privileges on each possible scope.	Privilege Tables for Hive and Impala

What's New in Cloudera Documentation in September, 2018

This section describes new topics added and major changes made to Cloudera documentation in September, 2018:

Table 50:

Product	What's New	Link
Apache Impala	Documented missing query options: <ul style="list-style-type: none"> • PARQUET_READ_STATISTICS • PARQUET_DICTIONARY_FILTERING 	<ul style="list-style-type: none"> • PARQUET_DICTIONARY_FILTERING Query Option • PARQUET_READ_STATISTICS Query Option • •
	Added a table of all Impala functions with links to each function as an alternative approach to having a page for each built-in function.	Impala Built-In Functions
	Reformatted the built-in functions docs format change for better readability.	
	Re-factored the Impala Authorization doc with the focus on Sentry privilege model and the de-emphasis on the policy file-based model.	Enabling Sentry Authorization for Impala
Apache Kudu	Added a best practice section in the Kudu-Spark Integration avoiding multiple Kudu clients per cluster.	Developing Applications With Apache Kudu
	Added the troubleshooting info on detecting ext2 and ext3 filesystems.	Troubleshooting Apache Kudu

Product	What's New	Link
Apache YARN	Added new topic that describes all aspects of creating and managing YARN ACLs.	Managing YARN ACLs

What's New in Cloudera Documentation in August, 2018

This section describes new topics added and major changes made to Cloudera documentation in August, 2018:

Table 51:

Product	What's New	Link
Cloudera Data Science Workbench	Added two new videos that demonstrate how to run experiments and deploy models with Cloudera Data Science Workbench.	<ul style="list-style-type: none"> • Experiments [YouTube] • Model Deployment [YouTube]
Apache Sentry	There is a new video on the Cloudera YouTube channel that shows how you can verify that your HDFS ACLs are synching with Sentry. The video also shows that URI privileges are not applied as ACLs in HDFS.	How to verify that HDFS ACLs are synching with Sentry
CDH - YARN	Updated the YARN tuning guide with new values.	Tuning YARN
Cloudera Altus	Description of Altus groups and their usage.	Groups
Cloudera Navigator	Added a new video that describes how to make sure your audit system is doing what you expect: are you collecting the right events? are you retaining them as long as you need them? are you archiving them where they are retrievable?	Navigator Audit Checkup Video [Youtube]
Workload Experience Manager (Workload XM)	Cloudera's Workload XM launched this month and with it a new documentation set that explains how to use this tool to gain in-depth understanding of the workloads you send to clusters managed by Cloudera Manager. It provides information that can be used for troubleshooting failed jobs and for optimizing slow jobs that run on those clusters.	Workload Experience Manager

What's New in Cloudera Documentation in July, 2018

This section describes new topics added and major changes made to Cloudera documentation in July, 2018:

Product	What's New	Link
Cloudera Data Science Workbench	Released Cloudera Data Science Workbench 1.4 with new features: Experiments and Models	1.4 Release Notes
	Reorganized documentation to align with major product components: Projects, Jobs, Experiments, Models, Engines, and Site Administration.	<ul style="list-style-type: none"> • Experiments • Model Deployment

Product	What's New	Link
	Improved LDAP/SAML experience with support for group filters.	LDAP and SAML
	New consolidated section for Engines in Cloudera Data Science Workbench.	Engines Overview
	New topic that describes how engines are used for experiments and models.	Engines for Experiments & Models
	This section also includes a topic that lists all the pre-installed packages in CDSW's Python and R kernels.	Pre-Installed Python and R Packages
	Provided more code samples that demonstrate how to access cluster data from CDSW.	Data Access
Cloudera Navigator	Added a new video that gives a light-hearted look at the Cloudera Navigator brand and helps identify the value of each of the Navigator components.	Navigator Brand Video [Youtube]
Reference Architectures	Cloudera Reference Architectures are now available in HTML format.	Reference Architectures
Apache Sentry	After upgrading to CDH 5.13.0 and above, some customers experience a period of time in which HDFS ACLs are not synched. Possible reasons for this problem are explained in the Release Notes, along with affected versions and fixes.	<ul style="list-style-type: none"> • After upgrade, HDFS ACLs are not synched due to Sentry taking too many snapshots • After upgrade, HDFS ACLs are not synched due to Sentry taking too long to retrieve the snapshot from the database
Cloudera Altus	Added a description of the optional ec2:DeleteKeyPair permission in the AWS cross-account role that determines how Altus generates key pairs for clusters.	Key Pair Permissions on EC2

What's New in Cloudera Documentation in June, 2018

This section describes new topics added and major changes made to Cloudera documentation in June, 2018:

Product	What's New	Link
Apache Hive - HiveServer2 High Availability	Added new command-line instructions for configuring a proxy load balancer to support HiveServer2 high availability on unmanaged clusters (those not managed by Cloudera Manager) with or without Kerberos.	Configuring HiveServer2 to Load Balance Behind a Proxy on Unmanaged Clusters
Apache Sentry	Clarified the GRANT ROLE statement on group name restrictions, such as character restrictions, how to use backticks with those restrictions, and OS group name requirements.	GRANT ROLE Statement
	Clarified the description of what happens to synchronized ACLs during Sentry service failure.	HDFS/Sentry Synchronized Permissions

Product	What's New	Link
	Added instructions for how to override Sentry's Kerberos prerequisite for the Hive metastore in Cloudera Manager.	Securing the Hive Metastore
	Added new Amazon S3 information on creating a table in a bucket.	Creating a Table in a Bucket
	Added new information on the privileges the Sentry Admin needs in HUE.	Hive SQL Syntax for Use with Sentry
	Added the SHOW CREATE VIEW operation to the Hive and Impala privilege tables.	Authorization Privilege Model for Hive and Impala
	Added a new example explaining how a user may see data from a database that they do not have access to if that data is in a view.	Authorization Privilege Model for Hive and Impala
CDK Powered by Apache Kafka	Updated Kafka Requirements and Supported Versions with additional information about compatibility: <ul style="list-style-type: none"> among client and broker versions between Spark and embedded Kafka clients. 	<ul style="list-style-type: none"> Kafka Client-Broker Compatibility across Kafka Versions Kafka Client Versions used by Apache Spark in CDH
Apache ZooKeeper	Provided instructions for configuring the ZooKeeper server for Kerberos authentication using Cloudera Manager.	Configuring ZooKeeper Server for Kerberos Authentication
Cloudera Manager	Added a new procedure that describes how to migrate from the Cloudera Manager Embedded PostgreSQL database server to an external PostgreSQL database.	Migrating from the Cloudera Manager Embedded PostgreSQL Database Server to an External PostgreSQL Database
Cloudera Navigator	Added information on lineage in Navigator: what information is collected and how it is used to create lineage diagrams, what entities are captured in the diagrams, and how diagrams change through the lifecycle of data assets.	Generating Lineage Diagrams
Cloudera Altus	Added a description of a new Altus environment option for secure clusters.	Enable Secure Clusters

What's New in Cloudera Documentation in May, 2018

This section describes new topics added and major changes made to Cloudera documentation in May, 2018:

Product	What's New	Link
Cloudera Upgrade	Added a new interactive topic that walks you through the steps to upgrade Cloudera Manager. You can select your operating system, upgrade version, and database type and a customized page displays the steps for your upgrade.	Upgrading Cloudera Manager Using Packages
	Added a new interactive topic that walks you through the steps to upgrade CDH using Cloudera Manager. You can select your Cloudera Manager version, CDH upgrade version, and other information	Upgrading CDH

Product	What's New	Link
	and a customized page displays the steps for your upgrade.	
HDFS Transparent Encryption	Extensively revised the KMS ACL topic, which now includes descriptions of all operations for each ACL class, as well as a diagram and explanation that guides readers through the process of how the KMS evaluates the various ACL classes.	Configuring KMS Access Control Lists (ACLs)
Key Trustee KMS HA	Added new documentation for a feature that provides logic to detect and warn users about a potential problem where the GPG private keys have not been properly synchronized across all Key Trustee KMS HA hosts.	<ul style="list-style-type: none"> • Validating Private Key Synchronization (Key Trustee KMS HA Only) • Enabling Key Trustee KMS High Availability • Key Trustee KMS Encryption Issues
Cloudera Navigator HSM KMS	Added a new topic to guide users through the steps to upgrade an HSM KMS.	Upgrading Cloudera Navigator HSM KMS
HBase	Added new content that describes how to configure and enable cell-level ACLs for HBase.	Configure Cell-Level Access Control Lists
Hue	Added new content that clarifies how to migrate the Hue database for MariaDB and MySQL.	MariaDB / MySQL
Cloudera Altus	Added information about defining custom tags for clusters. Added information about Altus support for CDH 5.14.	Creating a Cluster for AWS Creating a Cluster for Azure
	Restructured Altus documentation to create one doc set for Altus on AWS and Altus on Azure. Altus documentation now includes an Administration Guide and a Data Engineering Guide.	Overview of Cloudera Altus Overview of Altus Data Engineering

What's New in Cloudera Documentation in April, 2018

This section describes new topics added and major changes made to Cloudera documentation in April, 2018:

Product	What's New	Link
Cloudera Altus	Added a new topic that describes how to set up an Altus trial account.	Getting Started with a Trial Account

What's New in Cloudera Documentation in March, 2018

This section describes new topics added and major changes made to the Cloudera documentation library in March, 2018:

Product	What's New	Link
Cloudera Data Science Workbench	Added a new video that demonstrates how to get started with a Cloudera Data Science Workbench built-in template project.	CDSW Quickstart Demo [Youtube]
	Added new Known Issues for Cloudera Manager and CDH integration.	Known Issues

What's New in Cloudera Documentation

Product	What's New	Link
	Added a new topic on migrating a CDSW Deployment to Another Host.	Migrating a CDSW Deployment
	Revamped the Backup topic with detailed instructions.	Creating a Backup
	Added a new topic on how to uninstall Cloudera Data Science Workbench.	Uninstalling CDSW
JDK Requirements	Added new section on Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction requirements.	Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction
Navigator	Added a new example for Navigator Audit Server on how to use audit events to determine what caused a schema change to a table. Use audit reports to identify the user or process that may be causing unwanted changes.	Who ran which operation against a table?
Cloudera Director	Added a new topic on using custom DNS names and DNS servers with auto-TLS.	Using Custom DNS with Auto-TLS in AWS

What's New in Cloudera Documentation in February, 2018

This section describes new topics added and major changes made to Cloudera documentation in February, 2018:

Product	What's New	Link
Flume	The Apache Flume content is moved to a new Flume Guide. Information for configuring, using, and managing Flume is consolidated in the Flume Guide.	Flume Guide
HBase	The Apache HBase content is moved to a new HBase Guide. All the information for configuring, managing, and troubleshooting HBase is in one central location.	HBase Guide
Key HSM	Added a new section describing the file naming convention used for encryption zone keys.	Key Naming Convention
HDFS (Encryption)	Added a new section describing how to resolve an error that can occur when the KMS jute buffer size is insufficient to hold all the tokens.	KMS server jute buffer exception
Sentry	The Apache Sentry content is moved to a new Sentry Guide. The Sentry Guide contains information on configuring, using, and troubleshooting Sentry, as well as how-to guides.	Sentry Guide
Cloudera Altus	Added a new topic that describes how to use the Cloudera Altus SDK for Java.	Using the Altus SDK for Java

What's New in Cloudera Documentation in January, 2018

This section describes new topics added and major changes made to Cloudera documentation in January, 2018:

Product	What's New	Link
Cloudera Data Science Workbench	Released Cloudera Data Science Workbench 1.3.0.	<ul style="list-style-type: none">Cloudera Data Science Workbench 1.3.0 Release Notes

Product	What's New	Link
		<ul style="list-style-type: none"> • Known Issues and Limitations • New - Custom Project Templates • New - Rollback Cloudera Data Science Workbench • Updated Requirements and Supported Platforms • Updated CSD Upgrade Guide
Impala	Added tip about using Kudu Java API, instead of JDBC interface, for rapid insert operations.	Configuring Impala to Work with JDBC
	Added <code>DATE_TRUNC()</code> function.	Impala Date and Time Functions
	Added new upper limit for <code>BATCH_SIZE</code> query option.	BATCH_SIZE Query Option
	Added information about a new kind of runtime filter, the “min-max” filter, which applies to join queries involving Kudu tables.	Using Impala to Query Kudu Tables
	Added new conditional operators: <code>IS [NOT] TRUE</code> , <code>IS [NOT] FALSE</code> , and <code>IS [NOT] UNKNOWN</code> .	SQL Operators
	Added information about changes to the output of the <code>SET</code> statement, dividing the options into multiple groups, and hiding some groups by default. New <code>SET ALL</code> syntax shows all the option groups.	SET Statement
Added a new <code>impala-shell</code> option <code>--query_option</code> and configuration file section <code>[impala.query_options]</code> . These features both allow specifying values for query options when starting <code>impala-shell</code> .	impala-shell Configuration Options	
Kafka	Updated examples and removed deprecated properties for how to use Kafka with Flume.	Using Kafka with Flume
Kafka	Updated Kafka upgrade topic to include versions.	Rolling Upgrade to Kafka 3.0.x
Key Trustee KMS	Added new procedure for migrating from a Key Trustee KMS (KT KMS) to a Hardware Security Module KMS (HSM KMS).	Migrating from a Key Trustee KMS to an HSM KMS
Cloudera Manager	Added information on using Cloudera Manager to configure credentials for cluster access to Microsoft ADLS. This access is enabled for running Hive and Impala queries on tables backed by data stored in ADLS and to browse ADLS data using Hue.	Configuring ADLS Access Using Cloudera Manager
	Added information on how to enable performing minor maintenance on cluster hosts, Cloudera Manager now fully manages the host decommission and recommission process. You can specify whether or not to replicate under-replicated data blocks to other DataNodes to maintain the cluster's replication factor during a maintenance window.	Tuning and Troubleshooting Host Decommissioning
	Provided examples for how to use the API to manage BDR.	How To Automate BDR Replication with the Cloudera Manager API

Product	What's New	Link
	Added a video walkthrough for how to add a cluster to Cloudera Manager.	How to Add a Cluster to Cloudera Manager [YouTube]
Cloudera Director	Added information on Cloudera Director 2.7 configuration option to point to an organization's LDAP server so that users common credentials may be used to login to Cloudera Director. When enabling LDAP support, Cloudera Director's built in user management is disabled.	Configuring Cloudera Director Server for LDAP and Active Directory
	Added information that Cloudera Director can handle all aspects of Java installation on the instances that it allocates and configures for Cloudera Manager and CDH clusters, offering more flexibility while simplifying the process for users.	Deploying Java on Cluster Instances
	Added a configuration option to Cloudera Director's AWS plugin to accommodate regions like GovCloud and China, where EC2 cannot tag instances upon creation. The documentation now includes the procedure for configuring the plugin to use this option.	Configuring Tag-on-create for AWS GovCloud (US) and China (Beijing) Regions
Cloudera Navigator	Added information that metadata searches in Navigator now include the ability to group search results by common properties. Group by lets you use technical, managed, and custom metadata to quickly identify small files, active SQL users, table-creation trends, and other data aggregation trends revealed by metadata properties. The documentation includes some examples of how grouping search results can help you understand trends in your data and to find specific data assets.	Grouping Search Results Using Metadata
	Updated Navigator role names to more clearly reflect the privileges they provide. One specific change is that the privilege for editing the name and description metadata for Navigator entities is now part of the Managed & Custom Metadata Editor role. Users with that role or the Full Administrator role can add and update entity names and descriptions in the Navigator console.	Cloudera Navigator User Roles
	Added information that audit filtering now allows a "not like" operator.	Filtering Audit Events
	Added information on handling sensitive data that links to the Cloudera Manager log redaction details.	Sensitive Data
	The documentation now includes the specific metadata removed during Navigator Metadata Server purge tasks.	What Metadata is Purged?
Kudu	Added new features and updates to Kudu administration:	Kudu Administration

Product	What's New	Link
	<ul style="list-style-type: none"> You can now add data directories to an existing master or tablet server. Kudu tablet servers are resilient to disk failures that occur on a disk storing data blocks. The description of the workflow to migrate to multi-master cluster was improved with details and examples. The description of the workflow to recover from a dead Kudu master was improved with details. 	
	Specified how client applications connect to Kerberized Kudu servers.	Client Authentication to Secure Kudu Clusters
Cloudera Altus	Added a description of public keys for cluster creation.	Creating a Cluster for AWS

CDH 5 Release Notes

For links to the detailed change lists that describe the bug fixes and improvements to all of the CDH 5 projects, see [CDH Version and Packaging Information](#) on page 145.

For more information about installing and configuring CDH 5, see [Cloudera Installation Guide](#).

New Features and Changes in CDH 5



Note: There is no CDH 5.0.7, 5.1.1, 5.1.6, 5.2.2, 5.2.7, 5.3.7, or 5.4.6 release.

What's New in CDH 5.16.x

See below for new features in CDH 5.16.x, grouped by component:

Apache Hive

The following new features have been added to Apache Hive in CDH 5.16.1:

Fine Grained Privileges

Sentry and Hive introduced fine grained privileges to provide object-level privileges to roles.

Fine grained privileges adds the `CREATE` privilege, which allow users to create databases and tables. See the [Sentry Privileges](#) documentation for more information about the new privileges.

Object Ownership

Object ownership designates an owner for a database, table, or view in Sentry. The owner of an object has the equivalent of the `ALL` privilege on the object. See the [Object Ownership](#) documentation for information about enabling object ownership.

Because of the new object ownership feature, HMS stores the user that creates a table or database in Hive as the object owner. If object ownership is enabled, Sentry grants the user the `OWNER` privilege. Whether or not object ownership is enabled, HMS stores the user that creates the object as the object owner. Previously, HMS stored the `hive` user as the object owner.

The following statements were added to Hive to support object ownership via Sentry:

- `ALTER DATABASE SET OWNER`
- `ALTER TABLE SET OWNER`
- `SHOW GRANT USER`

Apache Impala

The following new features have been added to Apache Impala in CDH 5.16.1:

Fine Grained Privileges

Sentry and Impala introduced fine grained privileges to provide object-level privileges to roles.

Fine grained privileges include the `REFRESH` and `CREATE` privileges, which allow users to create databases and tables, and to execute commands that update metadata information on Impala databases and tables. See [Impala Sentry documentation](#) for the new privileges and the scopes of the objects that you can grant the new privileges on.

The following new privileges were added:

- The `REFRESH` privilege
- The `CREATE` privilege

- The `SELECT` and `INSERT` privileges on `SERVER`

If a role has `SELECT` or `INSERT` privilege on an object in Impala before upgrading to CDH 5.16.1, that role will automatically get the `REFRESH` privilege during the upgrade.

Object Ownership

Object ownership designates an owner for a database, table, or view in Sentry. The owner of an object has the `OWNER` privilege which is the equivalent of the `ALL` privilege on the object. See the [Object Ownership](#) documentation for information about enabling object ownership.

If the object ownership feature is enabled, Sentry grants the user the `OWNER` privilege. Whether or not object ownership is enabled, HMS stores the object creator as the default object owner. Previously, HMS stored the Kerberos user as the object owner.

The following statements were added to Impala to support object ownership via Sentry:

- `ALTER DATABASE SET OWNER`
- `ALTER TABLE SET OWNER`
- `ALTER VIEW SET OWNER`
- `SHOW GRANT USER`

Admission Control Enhancement

A new query option, [MAX MEM ESTIMATE FOR ADMISSION](#), was added. Use the new option to set an upper limit on the memory estimates of a query as a workaround for over-estimates precluding a query from being admitted.

General Performance Improvements

A new query option, [SHUFFLE DISTINCT EXPRS](#), controls the shuffling behavior when a query has both grouping and distinct expressions.

Metadata Performance Improvements

- Incremental Stats

The following enhancements improve Impala stability. The features reduce chances of having `catalogd` and `impalad` crash due to be out of memory when using incremental stats.

- Incremental stats are now compressed in memory in `catalogd`, reducing memory footprint in `catalogd`.
- Incremental stats are fetched on demand from `catalogd` by `impalad` coordinators. This enhancement reduces memory footprint for `impalad` coordinators and `statedored` and also reduces network requirements to broadcast metadata.

See [Loading Incremental Statistics from Catalogd](#) for details.

- Automatic Invalidation of Metadata



Note: This feature is experimental and not recommended for use in production clusters.

To keep the size of metadata bounded and to reduce the chances of `catalogd` cache running out of memory, this release introduces an automatic metadata invalidation feature with time-based and memory-based invalidation.

Automatic invalidation of metadata provides more stability with lower chances of running out of memory, but could potentially cause performance risks. The feature is turned off by default.

See [Startup Options for Automatic Invalidation of Metadata](#) for details.

Compatibility and Usability Enhancements

- Additional separators are supported between date and time in default `TIMESTAMP` format, specifically, the multi-space separator and the 'T' separator. See [TIMESTAMP Data Type](#) for more information on `TIMESTAMP` format.

CDH 5 Release Notes

- New hint placement is supported for `INSERT` statements. See [Optimizer Hints in Impala](#) for detail.
- The `REGEX_ESCAPE()` function was implemented for escaping special characters to treat them literally in string literals.
- `SHOW CREATE VIEW` was implemented with the same functionality as `SHOW CREATE TABLE`.
- Impala Shell can connect directly to `impalad` when configured with proxy load balancer and Kerberos. See [Configuration Options](#) for the new flag that enables the direct connection.

Apache Sentry

The following new features have been added to Apache Sentry in CDH 5.16.1:

Fine Grained Privileges

The `CREATE` and `REFRESH` (Impala only) privileges have been introduced to allow users to create databases, tables and functions, and to execute commands that update metadata information on Impala databases and tables.

For more information about the new privileges, see [Sentry Privileges](#).

Object Ownership

Object ownership designates an owner for a database, table, or view in Sentry. The owner of an object has the equivalent of the `ALL` privilege on the object.

In CDH 5.16.1, object ownership is disabled by default. For information about enabling and using object ownership, see [Object Ownership](#).

No Group Name Case Restrictions

Sentry no longer normalizes group name characters to be lowercase. Therefore, operating system group names do not need to be treated as case insensitive. In previous versions, Sentry modified capital letters in operating system group names to be lowercase.

OpenJDK Support

OpenJDK is now supported with Cloudera Manager and CDH 5.16.1.

For information about migrating your Oracle JDK to OpenJDK, see [Upgrading the JDK](#). To use OpenJDK in a new deployment, see [Install Java Development Kit](#).

Cloudera Bug: CDH-60334

Apache Kudu

The following new features have been added to Apache Hive in CDH 5.16.1:

Starting with Apache Kudu 1.5.0 / CDH 5.13.x, Kudu has been fully integrated into CDH. Kudu now ships as part of the CDH parcel and packages. For a complete list of new features and changes introduced in Kudu 1.7 / CDH 5.16.1, see [What's New in Apache Kudu](#) on page 1901.

Apache Phoenix

Apache Phoenix is available with CDH 5.16.2:

Starting with Apache Phoenix 4.14.1/ CDH 5.16.2, Phoenix parcel can be installed and used with CDH. You can now download and install the Phoenix parcel on CDH 5.16.2. For more information about installing and using Phoenix, see the [Phoenix Guide](#).

What's New in CDH 5.15.x

The following sections describe new features introduced in 5.15.0.

Apache Flume

You can use Cloudera Manager to configure Flume to communicate with Kafka sources, sinks, and channels over TLS. The **Kafka Service** property in Cloudera Manager allows you to select a dependent Kafka service for the Flume service. Cloudera Manager then creates `jaas.conf` and `flume.keytab` files and adds Kafka security properties to the Flume configuration file.



Note:

There is a known issue with this feature in 5.15.0, fixed in 5.15.1. For more information, see [Flume uses default values for Kafka TLS settings](#) on page 2522.

For more information about how to use Cloudera Manager to configure Flume and Kafka communication, see [Configuring Flume Security with Kafka](#).

Apache Hadoop

HDFS Immutable Snapshots

You can now enable immutable snapshots for HDFS. When you enable immutable snapshots, HDFS captures the sizes of files that are currently open (being written to). If immutable snapshots are not enabled, the sizes of those files in the snapshot will continue to grow based on any appends to the files. Therefore, immutable snapshots provide a more accurate point-in-time snapshot of HDFS.

This feature is useful for using snapshots with replication to prevent failures caused by file sizes growing in the source.

To enable immutable snapshots, open the **Cloudera Manager Admin Console**, select **<HDFS service> > Configuration** and search for the **Enable Immutable Snapshots** property.

You must enable immutable snapshots to use snapshot diff-based replication for Cloudera Manager's Backup and Disaster Recovery. For more information, see [What's New in Cloudera Manager 5.15.0](#) on page 2474.

Hue

End-user Data Catalog improvements:

- Simpler top table search
- Unification and caching of all SQL metadata (Hive, Navigator, NavOpt)

Apache Impala

- Added the `TABLESAMPLE` clause in the `COMPUTE STATS` statement. See [for the new TABLESAMPLE clause](#).
- Extended `COMPUTE STATS` to support a list of columns. See [for the new syntax](#).
- Added the new `COMPUTE_STATS_MIN_SAMPLE_SIZE` query option. The query option specifies the minimum number of bytes that will be scanned in `COMPUTE STATS TABLESAMPLE`, regardless of the user-supplied sampling percent. See [for information about the new option](#).
- Added a `TBLPROPERTY` for controlling stats extrapolation on a per-table basis: `impala.enable.stats.extrapolation=true/false`. See [for information about stats extrapolation](#).
- Added the new built-in `regex_escape` function. The function instructs Impala to interpret the following special characters literally rather than as special characters: `.\+*?[\^]${}!=!<>| :-`

See [for information about the regex_escape function](#).

- Enhanced the existing `ltrim` and `rtrim` string functions to accept an argument that specifies a set of characters to be trimmed from the input string. See [for information about the functions](#).
- Implemented `murmur_hash` function. See [for information about the new function](#).
- Introduced the support for the Kudu `DECIMAL` type in Kudu 1.7.0.
- Now Impala maps a signed integer logical type in Parquet to a supported Impala column type as below:
 - `INT_8` -> `TINYINT`
 - `INT_16` -> `SMALLINT`

- INT_32 -> INT
- INT_64 -> BIGINT
- Parquet dictionary filtering now works on nested data.
- Based on the existing Parquet column chunk level statistics `null_count`, Impala's Parquet scanner was enhanced to skip an entire row group if the `null_count` statistics indicate that all the values under the predicated column are `NULL` as no result rows would be returned from that row group.
- The Oracle-style hint placement for `INSERT` statements is now supported. See [for information on hints in Impala SQL](#).
- Insert plan hints for `CREATE TABLE AS SELECT` are now supported. See [for information on hints in Impala SQL](#).
- Improved concurrency of DDL and DML operations during catalog updates.
- The statestore update logic was improved to reduce issues, such as too many queries being admitted by different coordinators, or queries being queued for longer than necessary and blocking subsequent updates to different topics.
- The size limit for statestore updates was increased, and copying of the metadata and reduce the memory footprint were reduced. Now the catalog objects are passed and (de)compressed between FE and BE one at a time.

Apache Kudu

Starting with Apache Kudu 1.5.0 / CDH 5.13.x, Kudu has been fully integrated into CDH. Kudu now ships as part of the CDH parcel and packages. The documentation for Kudu has also been incorporated into the Cloudera Enterprise documentation [here](#).

For a complete list of new features and changes introduced in Kudu (in CDH 5.15), see [What's New in Apache Kudu](#) on page 1901.

Apache Spark

More flexibility to interpret `TIMESTAMP` values written by Impala. Setting the `spark.sql.parquet.int96TimestampConversion` configuration setting to `true` makes Spark interpret `TIMESTAMP` values, when reading from Parquet files written by Impala, without applying any adjustment from the UTC to the local time zone of the server. This behavior provides better interoperability for Parquet data written by Impala, which does not apply any time zone adjustment to `TIMESTAMP` values when reading or writing them.

What's New in CDH 5.14.x

The following sections describe new features introduced in 5.14.0.

Apache Impala

- **New Boolean operators:** `IS TRUE`, `IS NOT TRUE`, `IS FALSE`, `IS NOT FALSE`, `IS UNKNOWN`, and `IS NOT UNKNOWN`. These operators always return `TRUE` or `FALSE`, even if the left-hand expression evaluates to `NULL`. `IS [NOT] TRUE` and `IS [NOT] FALSE` are `NULL`-safe equivalents for equality and inequality comparisons. `IS [NOT] UNKNOWN` is identical to `IS [NOT] NULL`.
- The new built-in function `DATE_TRUNC()` lets you truncate a `TIMESTAMP` value to a given precision. For example, calling `date_trunc('hour', ts)` truncates `ts` to the beginning of the corresponding hour, with all minutes, seconds, milliseconds, and so on set to zero.
- The `impala-shell` command now accepts a command-line argument `--query_option=option=value` to set a query option for that shell session. You can specify multiple query options by supplying multiple instances of `--query_option` in the `impala-shell` invocation.

Configuration files for `impala-shell`, such as `.impalarc`, can also contain a section beginning with `[impala.query_options]`, followed by lines of the form `query_option=value`. These query options take effect by default for all `impala-shell` sessions that use the specified configuration file.

- The `SET` statement, when run without any argument, now divides its output into groups corresponding to “regular”, “advanced”, “development” (unfinished, not ready for general use), and “deprecated” (subject to removal). By

default, the “development” and “deprecated” option groups are not shown in `SET` output. Issue the statement `SET ALL` to see the options and values from all groups.

- Impala optimizes some join queries against Kudu tables by constructing runtime filters representing the minimum and maximum matching values from the join columns. The min-max information is “pushed down” to Kudu so that Kudu can prune unnecessary rows before sending the result set back to Impala.
- The allowed range for the `BATCH_SIZE` query option is now 0 to 65536. Formerly, the upper bound was undefined.
- Arithmetic expressions involving both `DECIMAL` and `FLOAT` or `DOUBLE` arguments now produce results of `DECIMAL` type instead of `DOUBLE`.

Apache Kudu

Starting with Apache Kudu 1.5.0 / CDH 5.13.x, Kudu has been fully integrated into CDH. Kudu now ships as part of the CDH parcel and packages. The documentation for Kudu has also been incorporated into the Cloudera Enterprise documentation [here](#).

For a complete list of new features and changes introduced in Kudu (in CDH 5.14), see [What's New in Apache Kudu](#) on page 1901.

Cloudera Search

In previous versions, the Lily HBase Indexer service only supported Sentry policy files for authorization. It now supports the Sentry Service. For more information, see [Configuring Lily HBase Indexer Security](#).

What's New in CDH 5.13.x

The following sections describe new features introduced in 5.13.0.

Apache Hive / Hive-on-Spark / HCatalog

- Support for dynamic partition pruning for map joins on Hive on Spark. Dynamic partition pruning (DPP) is a database optimization that can significantly decrease the amount of data that a query scans, thereby executing your workloads faster. It is disabled by default, but can be enabled by setting the `hive.spark.dynamic.partition.pruning.map.join.only` property to `true`. When enabled, DPP only triggers for queries where the join on the partitioned column is a map join. For details, see [Dynamic Partition Pruning for Hive Map Joins](#).
- Sentry supports Hive metastore high availability. In CDH 5.13 and above, you can use Sentry with Hive metastore high availability, with or without Sentry high availability. For information about the high availability architecture and steps to set up Sentry high availability, see [Sentry High Availability](#).
- Apache Pig now supports writing partitioned Hive tables in the Parquet format using HCatalog. For details, see [Using HCatalog to Write to Parquet Hive Tables with Pig](#).

Apache Impala

The following are some of the most significant new Impala features in this release:

- Improvements to memory management through the use of a buffer pool. This mechanism allows queries to use less memory, reserves the required memory during query startup, and reduces the frequency of out-of-memory errors. It makes query planning and memory estimation more accurate, so that if a query begins executing it is unlikely to encounter an out-of-memory error partway through. The memory buffer used during spill-to-disk processing is smaller: instead of 8 MiB, this buffer defaults to 2 MiB and Impala can reduce it to as little as 64 KiB when appropriate.

This feature includes new query options for fine-tuning memory areas used during query processing: `MIN_SPILLABLE_BUFFER_SIZE`, `DEFAULT_SPILLABLE_BUFFER_SIZE`, `MAX_ROW_SIZE`, and `BUFFER_POOL_LIMIT`.

- Improvements to the mechanism for caching HDFS file handles. This caching mechanism improves the performance and scalability of queries that access the same file multiple times, for example to retrieve different columns from a Parquet file. Caching the file handle across `open()` calls reduces the load on the HDFS NameNode.

This feature is currently disabled by default. It is enabled by setting a non-zero value for the `max_cached_file_handles` configuration setting. Currently, ETL processes that append to existing HDFS files or overwrite HDFS files in place can interact with this setting in a way that turns off short-circuit reads for some `impalad` hosts. See [HDFS-12528](#) for tracking information.

- A new command in `impala-shell`, `rerun` or its abbreviation `@`, lets you re-execute previous commands based on their numbering in the `history` output.
- You can specify the minimum required TLS/SSL version using the `--ssl_minimum_version` setting, for example `--ssl_minimum_version=tlsv1.2`.
- You can specify the set of allowed TLS ciphers using the `--ssl_cipher_list` configuration setting. See the output of `man ciphers` for the full set of keywords and notation allowed in the argument string.
- New or enhanced built-in functions:
 - `trunc()` can now apply to numeric types (`FLOAT`, `DOUBLE`, and `DECIMAL`) in addition to `TIMESTAMP`. Although this functionality was already available through the `truncate()` function, the new signatures for `trunc()` make it easier to port code from other popular database systems to Impala.
 - A new date/time function `utc_timestamp()` provides a simple way to get a stable, interoperable representation of a `TIMESTAMP` value without using a chain of functions to convert between representations and apply a specific timezone.
- The `CREATE TABLE LIKE PARQUET` statement can now handle Parquet files produced outside of Impala and containing `ENUM` types. The `ENUM` columns become `STRING` columns in the target table, and the `ENUM` values are turned into corresponding `STRING` values.
- Kudu enhancements:
 - You can now create a Kudu table without using a `PARTITION BY` clause. Kudu automatically creates a single partition to cover the entire possible range of values. This feature is intended for small lookup tables, where the overhead of partitioning does not make sense for tables that typically have a full table scan for each query.
 - More granular Sentry authorization for Kudu tables. Kudu tables can now use column-level privileges. The `SELECT` and `INSERT` statements can now use the corresponding `SELECT` and `INSERT` privileges. Other Kudu statements still require the `ALL` privilege.
 - The `ALTER TABLE` statement can modify a number of storage attributes for the columns of Kudu tables. You can use the `ALTER COLUMN` clause of `ALTER TABLE` along with the `SET` keyword to change the properties `DEFAULT`, `BLOCK_SIZE`, `ENCODING`, and `COMPRESSION`. You can use the `DROP DEFAULT` clause to remove the default value from a column.
- For non-Kudu tables, you can use the `ALTER TABLE` syntax `ALTER COLUMN col SET COMMENT 'text'` to change the comment for an individual column.

Apache Kudu

Starting with Apache Kudu 1.5.0 / CDH 5.13.x, Kudu has been fully integrated into CDH. Kudu now ships as part of the CDH parcel and packages. The documentation for Kudu has also been incorporated into the Cloudera Enterprise documentation [here](#).

For a complete list of new features and changes introduced in Kudu (in CDH 5.13), see [What's New in Apache Kudu](#) on page 1901.

Apache Pig

- Pig now supports writing partitioned Hive tables in the Parquet format using HCatalog. For details, see [Using HCatalog to Write to Parquet Hive Tables with Pig](#).

Apache Sentry

- High availability for Sentry provides automatic failover in the event that your primary Sentry host goes down or is unavailable. In CDH 5.13.0, you can have two Sentry hosts. For information about how to configure high availability for Sentry, see [Sentry High Availability](#).
- You must add Sentry to the list of groups that the Hive user can impersonate in the **Hive Metastore Access Control and Proxy User Groups Override** setting. See [Configuring the Sentry Service](#) for more information about Sentry service configuration.
- Sentry supports Hive metastore high availability.

What's New in CDH 5.12.x

The following sections describe new features introduced in 5.12.0.

Apache Hadoop

- Added documentation for using the HA nameservice of a remote cluster in the DistCp command line arguments. For more information, see [Using DistCp with Highly Available Remote Clusters](#).
- [HADOOP-14104](#): Added DistCp support for clusters with HDFS encryption zones.

Apache Hive / Hive-on-Spark

- Support for Microsoft Azure Data Lake Store (ADLS) as a secondary filesystem for both Hive on MapReduce2 (YARN) and Hive-on-Spark. You can now use both Hive on MapReduce2 and Hive-on-Spark to read and write data stored on ADLS.

See [Configuring ADLS Connectivity for CDH](#)

- The Hive `schematool` is integrated with Cloudera Manager where you can use it to upgrade or validate the Hive metastore schema.

See [Using the Hive Schema Tool in CDH](#) for details.

- [HIVE-1575](#): Added support for JSON arrays at the root level by the `get_json_object` function. For example:

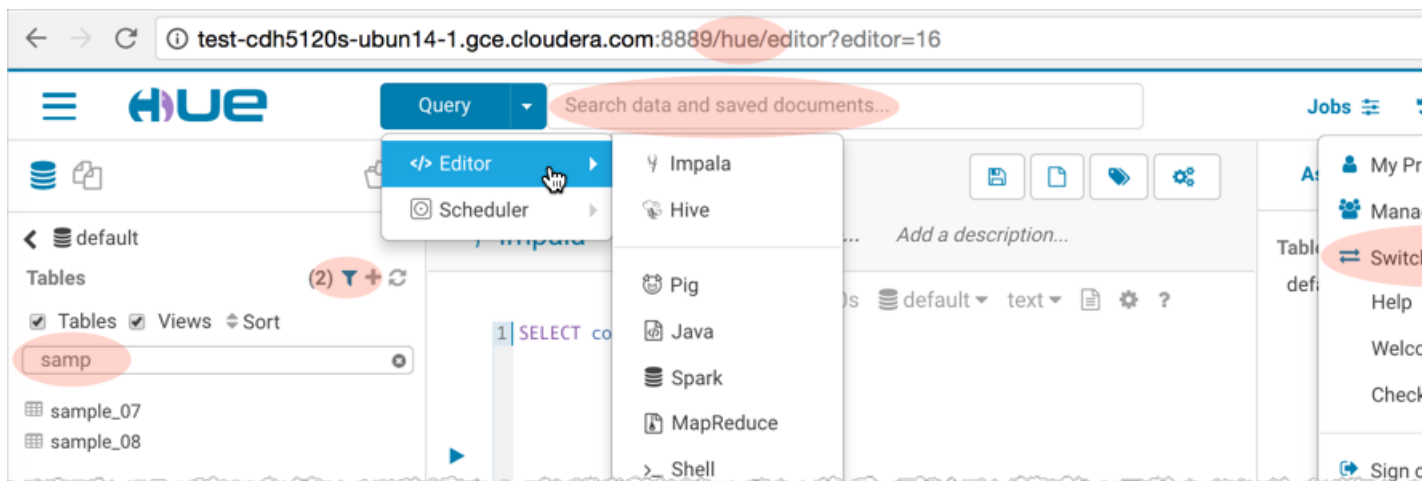
```
SELECT get_json_object('[1,2,3]', '$[0]')...
```

Hue

Hue 4 is out and jam-packed with great new features.

New Layout in Hue 4!

- Apps are consolidated under blue button—set your favorite as default landing page
- Top search bar lets you search for saved queries and other data
- Left and right assist panels let you search and filter schema objects
- Cursor position determines which of multiple queries to run
- New Pig editor, Job Designer, and Job Browser
- Access old Hue 3 layout under user drop down or remove "hue" from URL.



Load Balancer Added by Default

During a new installation of CDH/Hue, one Load Balancer is automatically promoted to ensure optimal performance—it can reduce the Hue server load by up to 90%! In existing clusters, administrators are prompted to add a load balancer role and users are then guided on how to enable it. See the Cloudera Blog on [Automatic HA](#).

Test LDAP Configuration

Verify your LDAP configuration, *on-the-fly*, with this new feature in Cloudera Manager under **Hue > Actions > Test LDAP Configuration**. See [Authenticate Hue Users with LDAP](#).

Navigator Optimizer Integrated (Phase 1)

With Navigator Optimizer enabled in Hue, popular tables, columns, joins, filters are displayed in the autocompleter. Risky statements, such as missing filters on partitioned tables, trigger an alert.

Navigator Search & Tag Enabled by Default

With Navigator enabled in Hue, you can search and tag metadata. This feature is now enabled by default (with Cloudera Navigator installed). See [How to Use Governance-Based Data Discovery](#).


Other Cool Features

- You can create partitioned tables from files
- Impala metadata is refreshed automatically
- SQL autocompleter handles more advanced corner cases
- Remote Load balancer works with SSL
- Query history is paginated!

Apache Impala

The following are some of the most significant new Impala features in this release:

- Impala can now read and write data stored on the Microsoft Azure Data Lake Store (ADLS). See [Using Impala with the Azure Data Lake Store \(ADLS\)](#) for details.
- New built-in functions:
 - A new string function, `replace()`, which is faster than `regexp_replace()` for simple string substitutions. See [Impala String Functions](#) for details.
 - A new conditional function, `nv12()`, which offers more flexibility than the `nv1()` function. It lets you return one value for `NOT NULL` arguments, and a different value for `NULL` arguments. See [Impala Conditional Functions](#) for details.
 - A new date/time function, `last_day()`, for finding the date corresponding to the last day of a particular month. See [Impala Date and Time Functions](#) for details.

- New syntax, `REFRESH FUNCTIONS db_name`, lets Impala recognize newly added functions, such as UDFs created through Hive. Impala scans the metadata for a specified database to locate the new functions, which is faster and more convenient than doing a full `INVALIDATE METADATA` operation.
 - The new `TABLESAMPLE` clause for the `SELECT` statement lets you specify a percentage of the overall data volume within a table to process during a query.
 - Startup flags for the `impalad` daemon, `is_executor` and `is_coordinator`, let you divide the work on a large, busy cluster between a small number of hosts acting as query coordinators, and a larger number of hosts acting as query executors. By default, each host can act in both roles, potentially introducing bottlenecks during heavily concurrent workloads. See [How to Configure Impala with Dedicated Coordinators](#) for details.
 - A new query option, `DEFAULT_JOIN_DISTRIBUTION_MODE`, lets you change the default assumption about how join queries should handle tables with no statistics. This can help to avoid out-of-memory conditions for join queries, without manual tuning to add the `/* +SHUFFLE */` hint for queries on large tables with missing statistics.
 - The `SORT BY` clause lets you create Parquet files with more efficient compression and smaller ranges of values for specified columns, allowing Impala to apply optimizations to skip reading data from Parquet files that do not contain any values that match equality and range operators in the `WHERE` clause. See [CREATE TABLE Statement](#) for details.
 - The `max_audit_event_log_files` lets you perform log rotation for the audit event log files, similar to the rotation for regular Impala log files.
 - Kudu enhancements:
 - The `ALTER TABLE` statement can specify more attributes for a Kudu table with the `ADD COLUMNS` clause. Now you can specify `[NOT] NULL`, `ENCODING COMPRESSION`, `DEFAULT`, and `BLOCK_SIZE`. See [ALTER TABLE Statement](#) for details.
 - The `TIMESTAMP` type is now available for Kudu tables.
-  **Note:** See [Handling Date, Time, or Timestamp Data with Kudu](#) for information about the tradeoffs between performance and convenience when using this data type. For high-performance applications, you might continue to use the `BIGINT` type to represent date/time values.
- The `INSERT` and `CREATE TABLE AS SELECT` statements are more efficient when writing to Kudu tables. Formerly, the overhead for the write operations could result in timeouts when writing large numbers of rows in a single operation.

Apache HBase

- Apache HBase now has ADLS support and recommendations for Azure deployment.
- Outside of cloud, HBase now has support for long-lived Spark applications via token renewal.

Apache Spark

Spark can read and write data on the Azure Data Lake Store (ADLS) cloud service. See [Accessing Data Stored in Azure Data Lake Store \(ADLS\) through Spark](#) for details.

S3Guard List-After-Delete Consistency

This release adds Delete Tracking for [S3Guard](#). Running CDH 5.11.0 clusters against the same DynamoDB table as clusters running CDH5.12.0 or higher is not supported. S3Guard in CDH5.11.0 does not work with the newer schema. CDH 5.11.1 and higher releases are compatible with the newer tables when running CDH 5.12.0 and higher.

S3Guard does not track deletes unless you are running a version with Delete Tracking support (CDH 5.12.0 and higher).

What's New in CDH 5.11.x

What's New in CDH 5.11.2

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.11.2](#) on page 2151.

What's New in CDH 5.11.1

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.11.1](#) on page 2154.

What's New In CDH 5.11.0

The following sections describe new features introduced in 5.11.0.

Apache Hadoop

- Supported Apache Tomcat TLS ciphers for HttpFS are configurable using the `HTTPFS_SSL_CIPHERS` environment variable.
- Supported Apache Tomcat TLS ciphers for the KMS are configurable using the `KMS_SSL_CIPHERS` environment variable.
- Amazon S3 Consistency with Metadata Caching (S3Guard)

Data written to Amazon S3 buckets is subject to the "eventual consistency" guarantee provided by Amazon Web Services (AWS), which means that data written to S3 may not be immediately available for queries and listing operations. This can cause failures in multi-step ETL workflows, where data from a previous step is not available to the next step. To mitigate these consistency issues you can now configure metadata caching for data stored in Amazon S3 using S3Guard. S3Guard requires that you provision a DynamoDB database from Amazon Web Services and configure S3Guard using the Cloudera Manager Admin Console or command-line tools. See [Configuring and Managing S3Guard](#).

- Amazon S3 Server-side Encryption with SSE-KMS

Clusters that use Amazon S3 storage can now use Amazon [Server-Side Encryption with AWS KMS–Managed Keys \(SSE-KMS\)](#) to encrypt data, so you now have two choices for data-at-rest encryption on Amazon S3 (SSE-S3, SSE-KMS). Use Cloudera Manager Admin Console to configure the cluster to use this new feature as detailed in [How to Configure Encryption for Amazon S3](#).

Apache HBase

- CDH 5.11.0 includes the introduction of “region server groups”. This feature allows for specific tables to be tied to specific region servers. The primary benefit of region server groups is application isolation. Multiple applications on HBase can be assured that there will be no i/o, cpu, and memory contention as long as they only access tables in mutually exclusive region server groups. There are trade-offs when using region server groups, notably performance. Tables that are assigned to region server groups will not be able to recruit the full hardware resources of the cluster, and the maximum possible throughput for tables assigned to region server groups composed of a subset of the region servers that comprise the cluster will be slower than those that are not.
- Use `MOB_COMPACT_PARTITION_POLICY` options to reduce the number of MOB files stored in HDFS. You can choose from daily, weekly, and monthly options.

Apache Hive

- Hive on Amazon S3 performance optimizations for:
 - [HIVE-14204](#) : Dynamic partitioning writes and the INSERT OVERWRITE statement
 - [HIVE-15546](#) : Parallel input path listing
 - [HIVE-13901](#), [HIVE_15879](#) : Performance and stability of the MSCK command for recovering partitions

See [Tuning Hive Performance on the Amazon S3 Filesystem](#)

- Support for Microsoft Azure Data Lake Store (ADLS) as a secondary filesystem for Hive on MapReduce2 (YARN). You can use Hive on MapReduce2 or Hive-on-Spark to read and write data stored on ADLS.

See [Configuring ADLS Connectivity for CDH](#).

- AWS cloud clusters can now share a single persistent instance of Amazon Relational Database Service (RDS) as the Hive metastore backend database, enabling persistent sharing of metadata beyond a cluster's life cycle.

See [How To Set Up a Shared Amazon RDS as Your Hive Metastore for CDH](#)

Hue

- Integrate Navigator with Hue: Phase 1, Metadata Discovery
 - Search and tag partitions, databases, views, tables, columns.
 - Off by default. Check both "Enable" fields in **Hue > Configuration > Cloudera Navigator**.
 - See [How to Use Governance-Based Data Discovery](#).
- Embed new create table wizard within Editor and Assist
 - Safely import multiple formats such as Kudu, Parquet, JSON, and CSV.
 - More easily create table partitions.
- Continued SQL improvements
 - Visually more pleasant colors and text.
 - No more hanging spinner in the Editor.
- [HUE-5742](#): Allow non-public PostgreSQL schemas.
- [HUE-5608](#): Add ability to DESC table without TABLE level privilege

Apache Impala

See [What's New in Apache Impala](#) on page 1869.

Apache Oozie

- Supported TLS ciphers for Apache Tomcat are configurable using the `OOZIE_HTTPS_CIPHERS` environment variable.

Apache Spark

Blacklisting. This feature reduces the chance of application failure, by not scheduling work on hosts that are experiencing intermittent disk failures. See [this blog post](#) for background information.

You can enable Kerberos authentication and TLS/SSL encryption for the Spark History Server through Cloudera Manager configuration settings, rather than including the password in clear text in an Advanced Configuration Snippet field. See these settings in the Cloudera Manager user interface:

- `history_server_spnego_enabled` - for Kerberos authentication
- `history_server_admin_users`
- `spark.ssl.historyServer.enabled`
- `spark.ssl.historyServer.protocol`
- `spark.ssl.historyServer.port`
- `spark.ssl.historyServer.enabledAlgorithms`
- `spark.ssl.historyServer.keyStore`
- `spark.ssl.historyServer.keyStorePassword`

With authentication enabled, only Kerberos-authorized users can read data from the Spark History Server, and non-admin users can only see information about their own jobs.

With TLS/SSL enabled, you provide the location of the keystore and its password, similar to the security configuration for other components.

Navigator lineage. The former Spark lineage extractor that was enabled through a safety valve is superseded by a more robust lineage collection mechanism. See [Apache Spark Known Issues](#) on page 2043 for some limitations and restrictions with this feature.

Support for Azure Data Lake Store (ADLS) as a secondary filesystem. You can use Spark jobs to read and write data stored on ADLS. Hive-on-Spark and Spark with Kudu are not currently supported for ADLS data.

Cloudera Search

- Supported TLS ciphers for Apache Tomcat are configurable using the `SOLR_CIPHERS_CONFIG` environment variable.

ZooKeeper

Server-Server Mutual Authentication

All ZooKeeper servers in an ensemble can now be configured to support quorum peer (server-server) mutual authentication, mitigating risk of spoofing by a rogue server on an unsecured network. The feature leverages Kerberos authentication through the SASL framework, so Kerberos is required.

This feature is easy to enable using Cloudera Manager Admin Console. See [Enabling Server-Server Mutual Authentication](#) in the [ZooKeeper Authentication](#) page of [Cloudera Security](#) for details.

What's New in CDH 5.10.x

What's New in CDH 5.10.2

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.10.2](#) on page 2168.

What's New in CDH 5.10.1

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.10.1](#) on page 2172.

What's New In CDH 5.10.0

Apache HBase

- Replication debugging functionality has been added.
- G1 Garbage Collection performance has been improved.

Apache Hive

- [HIVE-14269](#): Enhanced write performance for Hive tables stored on Amazon S3. See [Tuning Apache Hive Performance on the Amazon S3 Filesystem in CDH](#).

Hue

- [HUE-4560](#): Auto-configure Hue connectivity to S3 for all components. See [How to Enable S3 Cloud Storage in Hue](#).

Apache Impala

See [New Features in Impala 2.8.x / CDH 5.10.x](#) on page 1875.

Apache Oozie

[OOZIE-2225](#) In addition to the filter parameters `name`, `user`, `group`, `status`, `frequency` and `unit`, Oozie now supports a wild card parameter that matches a partial name, a partial user, or a complete ID.

For example, a query with `text= 'asdf'` would return all jobs that satisfy ANY of the following conditions.

1. `job.name.contains(text)`
2. `job.user.contains(text)`
3. `job.id = text`

Cloudera Search

- The `solrctl` command can be run in debug mode using the `--debug` or `--trace` parameters.

YARN

- **Expose maxResources per user in YARN Dynamic Resource Pools**

YARN Dynamic Resource Pools now supports default capacity limits (**Max Resources**) that automatically apply to all child pools of any resource pool. You can control the YARN resources available to any user or group on the cluster by configuring these settings on a parent pool, and then use placement rules to auto-create child pools per user or group.

- **UX improvements in the Pools page**

Improvements include:

- Every input row is now clickable and takes the user to that input field in the edit dialog.
- Split vcores (min/max) and memory (min/max) into min resources and max resources.
- Renamed **Configuration Sets** tab to **Resource Limits**.

What's New In CDH 5.9.x

What's New in CDH 5.9.3

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.9.3](#) on page 2194.

What's New in CDH 5.9.2

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.9.2](#) on page 2197.

What's New in CDH 5.9.1

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.9.1](#) on page 2203.

What's New in CDH 5.9.0

Apache Hadoop

- CDH 5.9 supports temporary credentials obtained from AWS Security Token Service (STS) to log in to Amazon S3.

Apache HBase

- A tool has been added--`org.apache.hadoop.hbase.replication.regionserver.DumpReplicationQueues`--to dump existing replication peers, configurations, and queues when using HBase replication. The tool includes two flags:
 - `--distributed` - Polls each replication server for information about the replication queues being processed on this replication server. By default, this is not enabled, and the information about the replication queues and configuration is obtained from ZooKeeper.
 - `--hdfs` When `--distributed` is used, this flag attempts to calculate the total size of the WAL files used by the replication queues. Because multiple peers can be configured, this value can be overestimated.

For more information, see [Class DumpReplicationQueues](#).

- Metrics have been added that expose the amount of replayed work occurring in the HBase replication system. For more information on these metrics, see [Replication Metrics](#) in the *Apache HBase Reference Guide*.

Apache Hive

- [HIVE-14270](#) : Added parameters to optimize write performance for Hive tables and partitions that are stored on Amazon S3. See [Optimizing Hive Write Performance on Amazon S3](#).

Hue

- [HUE-2915](#): Integrates Hue with [Amazon S3](#). You can now access both S3 and HDFS in the File Browser, create tables from files in S3, and save query results in S3. See how to [Enable S3 Cloud Storage](#).
- [HUE-4039](#): Improves SQL Autocompleter. The new Autocompleter deeply understands Hive and Impala SQL dialects and provides smart suggestions based on statement structure and cursor position. See how to manually [Enable and Disable Autocompleter](#).
- [HUE-3877](#): Adds support for Amazon RDS. You can now deploy Hue against an Amazon RDS database instance with MySQL, PostgreSQL, and Oracle engines.
- Rebase of Hue on upstream [Hue 3.11](#).

Apache Impala

- Performance improvements:
 - [\[IMPALA-3206\]](#) Speedup for queries against `DECIMAL` columns in Avro tables. The code that parses `DECIMAL` values from Avro now uses native code generation.
 - [\[IMPALA-3674\]](#) Improved efficiency in LLVM code generation can reduce codegen time, especially for short queries.
 - [\[IMPALA-2979\]](#) Improvements to scheduling on worker nodes, enabled by the `REPLICA_PREFERENCE` query option. See [REPLICA_PREFERENCE Query Option \(CDH 5.9 or higher only\)](#) for details.
- [\[IMPALA-1683\]](#) The `REFRESH` statement can be applied to a single partition, rather than the entire table. See [REFRESH Statement](#) and [Refreshing a Single Partition](#) for details.
- Improvements to the Impala web user interface:
 - [\[IMPALA-2767\]](#) You can now force a session to expire by clicking a link in the web UI, on the `/sessions` tab.
 - [\[IMPALA-3715\]](#) The `/memz` tab includes more information about Impala memory usage.
 - [\[IMPALA-3716\]](#) The **Details** page for a query now includes a **Memory** tab.
- [\[IMPALA-3499\]](#) Scalability improvements to the catalog server. Impala handles internal communication more efficiently for tables with large numbers of columns and partitions, where the size of the metadata exceeds 2 GiB.
- [\[IMPALA-3677\]](#) You can send a `SIGUSR1` signal to any Impala-related daemon to write a Breakpad minidump. For advanced troubleshooting, you can now produce a minidump without triggering a crash. See [Breakpad Minidumps for Impala \(CDH 5.8 or higher only\)](#) for details about the Breakpad minidump feature.
- [\[IMPALA-3687\]](#) The schema reconciliation rules for Avro tables have changed slightly for `CHAR` and `VARCHAR` columns. Now, if the definition of such a column is changed in the Avro schema file, the column retains its `CHAR` or `VARCHAR` type as specified in the SQL definition, but the column name and comment from the Avro schema file take precedence. See [Creating Avro Tables](#) for details about column definitions in Avro tables.
- [\[IMPALA-3575\]](#) Some network operations now have additional timeout and retry settings. The extra configuration helps avoid failed queries for transient network problems, to avoid hangs when a sender or receiver fails in the middle of a network transmission, and to make cancellation requests more reliable despite network issues.

Apache Oozie

Oozie adds a new database tool for migration and upgrade from Apache Derby (or any other supported database). For more information, see [How to Use the New Apache Oozie Migration Tool](#).

Apache Sentry

- Sentry supports data on Amazon RDS and can secure URIs with an RDS schema.
- [SENTRY-1233](#) - Logging improvements for `SentryConfigToolSolr`.

- [SENTRY-1119](#) - Allow data engines to obtain the `ActionFactory` directly from the configuration, instead of having hardcoded component-specific classes. This will allow external data engines to integrate with Sentry easily.
- [SENTRY-1229](#) - Added a basic configurable cache to `SentryGenericProviderBackend`.

Apache Spark

- You can now set up AWS credentials for Spark with the Hadoop credential provider, to avoid exposing the AWS secret key in configuration files.

Apache Sqoop

- The mainframe import module extension has been added to support data sets on tape.

Cloudera Search

- The Solr watchdog is now configured to use the fully qualified domain name (FQDN) of the host on which the Solr process is running (instead of `127.0.0.1`). You can override this configuration by setting `SOLR_HOSTNAME` environment variable to appropriate value (before starting the Solr server).
- Cloudera Search adds support for index snapshots. For more information on how to back up, migrate, or restore your indexed data, see [Backing Up and Restoring Cloudera Search](#).

What's New In CDH 5.8.x

What's New in CDH 5.8.5

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.8.5](#) on page 2232.

What's New in CDH 5.8.4

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.8.4](#) on page 2247.

What's New in CDH 5.8.3

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.8.3](#) on page 2251.

What's New in CDH 5.8.2

- Support for HDFS Disk Balancer. For details, see <http://hadoop.apache.org/docs/r3.0.0-alpha1/hadoop-project-dist/hadoop-hdfs/HDFSDiskbalancer.html>.

What's New in CDH 5.8.0

The following sections describe new features introduced in CDH 5.8.0.

Operating System Support

- **Operating Systems** - Support for Debian 8.2.

Apache HBase

- Additional metrics have been added to monitor garbage collection pauses and other external pauses that might cause a server process to momentarily block any request processing.
- New throughput and Bloom filter metrics have been added to the existing HBase microbenchmarks.
- Downstream users of CDH who build on top of Apache HBase can now pull in fewer transitive dependencies by relying on shaded client artifacts. To do so, use the maven artifact `hbase-shaded-client` where you would normally use the `hbase-client` artifact. All HBase APIs remain the same.



Note: In a previous release of CDH, this maven artifact exists but does not contain any of the needed classes to interact with an HBase cluster.

- The HMaster Web UI now shows the aggregate of all space being used by snapshots.

Hue

General Features:

- Rebase of Hue on upstream [Hue 3.10](#).
- Refactor of Hue Infrastructure:
 - Performance is further optimized for large numbers of databases and tables.
 - Exporting and importing documents is improved.

SQL Editor and Browser:

- Revamp of the Hue SQL Application:
 - Editor is redesigned to create a single-page experience.
 - Code editor is redesigned with enhanced auto-complete, keyboard shortcuts, search and replace, and more.
 - A live status of the query history displays as an icon



Admin and Security:

- [HUE-3386](#): Users are auto-logged out when TTL Expires. See [Securing Sessions](#).
- [Hue-3808](#): Users can do live DEBUG log toggling.
- A new step in the Cloudera Manager Add Service wizard helps configure and test an external Hue database.

Search:

- [SENTRY-1217](#): Users can grant Sentry Solr privileges.

Oozie:

- [Hue-3464](#): The dashboard and editor are decoupled for granular access.
- Saved Hive queries can be dragged and dropped into a workflow.

Apache Impala

See [New Features in Impala 2.6.x / CDH 5.8.x](#) on page 1878.

Apache Oozie

- [OOZIE-2330](#) : The Spark Action now allows <file> and <archive> elements. It also omits <job-tracker> and <name-node> elements to allow for the use of the global or default values.
- The launcher job no longer uses YARN's uber mode by default.

Cloudera Search

- Cloudera Search adds support for storing permissions in the Sentry service. You can enable storing permissions in the Sentry service by [Enabling the Sentry Service for Solr](#). If you have already configured Sentry's policy file-based approach, you can migrate existing authorization settings as described in [Migrating from Sentry Policy Files to the Sentry Service](#). `solrctl` has been extended to support:
 - Migrating existing policy files to the Sentry service
 - Managing managing permissions in the Sentry service

Apache Sentry

- Sentry supports data stored on Amazon S3 and can secure URIs with an S3 schema.
- Cloudera Search adds support for storing permissions in the Sentry service. You can enable storing permissions in the Sentry service by [Enabling the Sentry Service for Solr](#). If you have already configured Sentry's policy file-based approach, you can migrate existing authorization settings as described in [Migrating from Sentry Policy Files to the Sentry Service](#). `solrctl` has been extended to support:
 - Migrating existing policy files to the Sentry service
 - Managing managing permissions in the Sentry service

- [SENTRY-1175](#): Improved usability for Sentry URIs and URI privileges. If URIs in Hive DDL statements or URI privileges lack scheme and authority components, Sentry automatically completes such URIs by applying the default scheme and authority based on the HDFS configuration provided to Sentry.
- **Performance Improvements**
 - [SENTRY-1293](#): ResourceAuthorizationProvider.doHasAccess no longer performs expensive operations to convert string permissions to Privilege objects.
 - [SENTRY-1292](#): Reordered the DBModelAction EnumSet to improve authorization performance.

Key Trustee Server

Auto backup on install of Key Trustee Server

Key Trustee server backs up automatically on first run when using parcels and Cloudera Manager, and sets up a cron job (hourly) for continuous backup.

Auto backup on install of KMS

Key Trustee KMS backs up automatically on first run when using parcels and Cloudera Manager.

What's New In CDH 5.7.x

What's New in CDH 5.7.0

The following sections describe new features introduced in CDH 5.7.0.

Operating System Support

- **Operating Systems** - Support for:
 - RHEL/CentOS 6.6, 6.7, 7.1, and 7.2
 - Oracle Enterprise Linux (OEL) 7.1 and 7.2
 - SUSE Linux Enterprise Server (SLES) 11 with Service Packs 2, 3, 4
 - Debian: Wheezy 7.0, 7.1, and 7.8



Important: Cloudera supports RHEL 7 with the following limitations:

- Only RHEL 7.3, 7.2 and 7.1 are supported. RHEL 7.0 is not supported.
- RHEL 7.1 is only supported with CDH 5.5 and higher.
- RHEL 7.2 is only supported with CDH 5.7 and higher.
- Only new installations of RHEL 7.2 and 7.1 are supported by Cloudera. For upgrades to RHEL 7.1 or 7.2, contact your OS vendor and see [Does Red Hat support upgrades between major versions of Red Hat Enterprise Linux?](#)

Apache Hadoop

- Improve support for heterogeneous storage. [Heterogeneous Storage Management](#) (HSM) is now supported natively in Cloudera Manager.
- [HADOOP-10651](#) - Service access can be restricted by IP and hostname. You can now define a whitelist and blacklist per service. The default whitelist is * and the default blacklist is empty.
- [HADOOP-12764](#) - You can configure `KMS maxHttpHeaderSize`. The default value of `KMS maxHttpHeaderSize` increased from 4096 to 65536 and is now configurable in `service.xml`.
- [HDFS-7279](#) - In CDH 5.5.0 and higher, DataNode WebHDFS implementation uses Netty as an HTTP server instead of Jetty. With improved buffer and connection management, Netty lowers the risk for DataNode latency and OutOfMemoryError (OOM).
- [HDFS-8873](#) - You can rate-limit the directoryScanner. A new configuration property, `dfs.datanode.directoryscan.throttle.limit.ms.per.sec`, allows you to reduce the impact on disk performance of directory scanning. The default value is 1000.

- [HDFS-9260](#) - Garbage collection of full block reports is improved. Data structures for `BlockInfo` and replicas were changed to keep them sorted. This allows for faster and easier garbage collection of full block reports.

Apache HBase

See also [Apache HBase Known Issues](#) on page 1970.

- CDH 5.7.0 adds support for a snapshot owner for a table. To configure a table snapshot owner, set the `OWNER` attribute on the table to a valid HBase user. By default, the table owner is the user who created the table. The table owner and the table creator can restore the table from a snapshot.
- You can set an HDFS storage policy to store write-ahead logs (WALs) on solid-state drives (SSDs) or a mix of SSDs and spinning disks.
- Configuration for snapshot timeouts has been simplified. A single configuration option, `hbase.snapshot.master.timeout.millis`, controls how long the master waits for a response from the `RegionServer` before timing out. The default timeout value has been increased from 60,000 milliseconds to 300,000 to accommodate larger tables.
- You can optionally balance a table's regions by size by setting the `hbase.normalizer.enabled` property to `true`. The default value is `false`. To configure the time interval for the HMaster to check for region balance, set the `hbase.normalizer.period` property to the interval, in milliseconds. The default value is 1800000, or 30 minutes.
- You can configure parallel request cancellation for multi-get operations using the `hbase.client.replica.interrupt.multiget` configuration property, using an advanced configuration snippet in Cloudera Manager, or in `hbase-site.xml` if you do not use Cloudera Manager.
- The `hbase-spark` module has been added, which provides support for using HBase data in Spark jobs using `HBaseContext` and `JavaHBaseContext` contexts. See the [HBase and Spark chapter](#) of the Apache HBase Reference Guide for details about building a Spark application with HBase support.
- The REST API now supports creating, reading, updating, and deleting namespaces.
- If you use the G1 garbage collector, you can disable the `BoundedByteBufferPool`.
- The HBase web UI includes graphical tools for managing MemStore and StoreFile details for a region.
- The HBase web UI displays the number of regions a table has.
- Two new methods of the Scan API, `setColumnFamilyTimeRange` and `getColumnFamilyTimeRange`, allow you to limit Scan results to versions of columns within a specified timestamp range.
- A new API, `MultiRowRangeFilter`, allows you to scan multiple row ranges in a single scan.
- A new client API, `SecurityCapability`, enables you to check whether the HBase server supports cell-level security.
- When using the `scan` command in HBase Shell, you can use the `ROWPREFIXFILTER` option to include only rows matching a given prefix, in addition to any other filters you specify. For example:

```
hbase> scan 't1', {ROWPREFIXFILTER => 'row2',
                  FILTER => (QualifierFilter (>=, 'binary:xyz')) AND
                  (TimestampsFilter ( 123, 456))}
```

- The new `get_splits` HBase Shell command returns the split points for a table. For example:

```
hbase> get_splits 't2'
Total number of splits = 5
=> ["", "10", "20", "30", "40"]
```

- Three new commands relating to the region normalizer have been added to the HBase Shell:
 - `normalizer_enabled` checks whether the region normalizer is enabled.
 - `normalizer_switch` toggles the region normalizer on or off.
 - `normalize` runs the region normalizer if it is enabled.

- When configuring a ZooKeeper quorum for HBase, you can now specify the port for each ZooKeeper host separately, instead of using the same port for each ZooKeeper host. The configuration property `hbase.zookeeper.clientPort` is no longer required. For example:

```
<property>
<name>hbase.zookeeper.quorum</name>
<value>zk1.example.com:2181,zk2.example.com:20000,zk3.example.com:31111</value>
</property>
```

- A new configuration option, `hbase.regionserver.hostname`, was added, but Cloudera recommends against its use. See [Compatibility Notes for CDH 5.7](#) on page 1922.
- Two new configuration options allow you to disable loading of all coprocessors or only table-level coprocessors: `hbase.coprocessor.enabled` (which defaults to `true`) and `hbase.coprocessor.user.enabled` (which also defaults to `true`). Cloudera does not recommend disabling HBase-wide coprocessors, because security functionality is implemented using coprocessors. However, disabling table-level coprocessors may be appropriate in some cases.
- A new configuration option, `hbase.loadincremental.validate.hfile`, allows you to skip validation of HFiles during a bulk load operation when set to `false`. The default setting is `true`.
- The default `PermSize` for HBase processes is now set to 128 MB. This setting is effective only for JDK 7, because JDK 8 ignores the `PermSize` setting. To change this setting, edit the **HBase Client Environment Advanced Configuration Snippet (Safety Valve) for `hbase-env.sh`** if you use Cloudera Manager or `conf/hbase-env.sh` otherwise.
- In CDH 5.6 and lower, the `HBaseFsck#checkRegionConsistency()` method would throw an `IOException` if a region repair operation timed out after `hbase.hbck.assign.timeout` (which defaults to 120 seconds). This exception would the entire `hbck` operation to fail. In CDH 5.7.0, if the region being repaired is not `hbase:meta` or another system table, the region is skipped, an error is logged, and the `hbck` operation continues. This new behavior is disabled by default; to enable it, set the `hbase.hbck.skipped.regions.limit` option to an integer greater than 0. If more than this number of regions is skipped, the `hbck` operation fails.
- Two new options, `-exclusive` and `-disableBalancer`, have been added to the `hbck` utility. The `hbck` utility now runs without locks unless in fixer mode, and the balancer is only disabled in fixer mode, by default. You can disable these options to retain the old behavior, but Cloudera recommends using the new default behavior.
- Two new MapReduce jobs, `SyncTable` and `HashTable`, allow you to synchronize two different HBase tables that are each receiving live writes. To print usage instructions, run the job with no arguments. These examples show how to run these jobs:

```
$ bin/hbase org.apache.hadoop.hbase.mapreduce.SyncTable \
--dryrun=true \
--sourcezkcluster=zk1.example.com,zk2.example.com,zk3.example.com:2181:/hbase \
hdfs://nn:9000/hashes/tableA \
tableA tableA
```

```
$ bin/hbase org.apache.hadoop.hbase.mapreduce.HashTable \
--batchsize=32000 \
--numhashfiles=50 \
--starttime=1265875194289 \
--endtime=1265878794289 \
--families=cf2,cf3 \
TestTable /hashes/testTable
```

- The `CopyTable` command now allows you to override the `org.apache.hadoop.hbase.mapreduce.TableOutputFormat` property by prefixing the property keys with the `hbase.mapred.output.` prefix. For example, `hbase.mapred.output.hbase.security.authentication` is passed to `CopyTable` as `hbase.security.authentication`. This is useful when directing output to a peer cluster with different configuration settings.
- Several improvements have been made to the HBase canary, including:
 - Sniffing of regions and RegionServers is now parallel to improve performance in large clusters with more than 1000 regions and more than 500 RegionServers.

- The canary sets `cachecheck` to `false` when performing Gets and Scans to avoid influencing the BlockCache.
- `FirstKeyOnlyFilter` is used during Gets and Scans to improve performance in a flat wide table.
- A region is selected at random when sniffing a RegionServer.
- The sink class used by the canary is now configurable.
- A new flag, `-allRegions`, sniffs all regions on a RegionServer if running in RegionServer mode.
- Distributed log replay has been disabled in CDH 5.7 and higher, due to reliability issues. Distributed log replay is unsupported.
- A new configuration option, `hbase.hfile.drop.behind.compaction`, causes the OS-level filesystem cache to be dropped behind compactions. This provides significant performance improvements on large compactions. To disable this behavior, set the option to `false`. It defaults to `true`.
- Two new configuration options, `hbase.hstore.compaction.max.size` and `hbase.hstore.compaction.max.size.offpeak`, allow you to specify a maximum compaction size during normal hours and during off-peak hours, if the off-peak feature is used. If unspecified, `hbase.hstore.compaction.max.size` defaults to 9,223,372,036,854,775,807 (`Long.MAX_VALUE`), which is essentially unbounded.
- The HDFS replication factor can now be specified per column family by setting the `DFS_REPLICATION` attribute of the column family to the desired number of replicas, either at table creation or by altering an existing table schema. If set to 0, the default replication factor is used. If fewer than the desired number of replicas exist, the HDFS FSCK utility reports it.
- If you use the ChaosMonkey tool, you can low-load a custom ChaosMonkey implementation by passing the class to the `-m` or `--monkey` option of the ChaosMonkey tool, in the same way that you would normally pass `SLOW` or `CALM`.
- A new configuration option, `hbase.use.dynamic.jars`, allows you to disable the dynamic classloader if set to `false`.

Apache Hive

- [HIVE-9298](#) - Supports reading alternate timestamp formats. The SerDe property, `timestamp.formats`, is added to allow you to pass in a comma-delimited list of alternate timestamp formats. For example, the following ALTER TABLE statement (in this case, with Joda date-time parsing) adds: `yyyy-MM-dd'T'HH:mm:ss` and milliseconds since [Unix epoch](#), represented by the special case pattern `millis`.

```
ALTER TABLE timestamp_formats SET SERDEPROPERTIES
 ("timestamp.formats"="yyyy-MM-dd'T'HH:mm:ss,millis");
```

- [HIVE-7292](#) - Hive on Spark (HoS) General Availability. This release supports running Hive jobs using Spark as an additional execution backend. This improves performance for Hive queries, especially those involving multiple reduce stages.
- [HIVE-12338](#) - Add Web UI to HiveServer2. Exposes a Web UI on the HiveServer2 service to surface process metrics (JVM stats, open session, and similar) and per-query runtime information (query plan, performance logs). The Web UI is available on port 10002 by default.
- [HIVE-10115](#) - Handle LDAP and Kerberos authentication on the same HiveServer2 interface. In a Kerberized cluster when alternate authentication is enabled on HiveServer2, it accepts Kerberos authentication. Before this enhancement, when LDAP authentication to HiveServer2 was enabled, the service blocked acceptance of other authentication mechanisms.
- Hive Metastore / HiveServer2 metadata scalability improvements. Fixes to improve scalability and performance include support for DirectSQL by default for metastore operations and incremental memory usage improvements.
- [HIVE-12271](#) - Add additional catalog and query execution metrics for Hive Metastore / HiveServer2 and integrate with Cloudera Manager. Additional Hive metrics include catalog information (number of databases, tables, and partitions in Hive Metastore) and query execution stats (number of worker threads used, job submission times, and statistics for planning and compilation times). All metrics are available in Cloudera Manager.

Hue

- Hive metastore service improvements make browsing Hive data faster and easier.
- SQL editor improvements:

- SQL assist scales to thousands of tables and databases.
- Improved row and column headers.
- Button to delete the query history.
- Button to format SQL.
- Security improvements:
 - Timeout property for idle sessions.
 - Customizable splash screen on the login page.
 - Support for password-protected SAML certificates.
 - Custom xmlsec1 binary for multiple SAML protocols.
 - Ability to synchronize groups on login for any authentication backend.
- Oozie application improvements:
 - Display graph of external workflow.
 - Option to dry run jobs on submission.
 - User timezone recognition.
 - Automatic email on failure.
 - Ability to execute individual actions.

Apache Impala

See [New Features in Impala 2.5.x / CDH 5.7.x](#) on page 1880.

MapReduce

- The `mapred job -history` command provides an efficient way to fetch history for MapReduce jobs.

Apache Oozie

- [OOZIE-2411](#) - The Oozie email action now supports `BCC:`, as well as `TO:` and `CC:`.

Cloudera Search

- Improvements to `loadSolr`. For additional information about these improvements, see the [Morphlines Reference Guide](#). Improvements include:
 - `loadSolr` retries SolrJ requests if exceptions, such as connection resets, occur. This is useful in cases such as temporary Solr server overloads or transient connectivity failures. This behavior is enabled by default, but can be disabled by setting the Java system property `org.kitesdk.morphline.solr.LoadSolrBuilder.disableRetryPolicyByDefault=true`.
 - `loadSolr` includes an optional rate-limiting parameter that is used to set the maximum number of morphline records to ingest per second. The default value is no limit.

Apache Spark

- Spark is rebased on Apache Spark 1.6.0.
- [SPARK-10000](#) - Spark 1.6.0 includes a new unified memory manager. The new memory manager is turned off by default (unlike Apache Spark 1.6.0), to make it easier for users to migrate existing workloads, but it is supported.
- [SPARK-9999](#) - Spark 1.6.0 introduces a new Dataset API. However this API is experimental, likely to undergo some changes, and unsupported.
- [SPARK-6028](#) and [SPARK-6230](#) - Encryption support for Spark RPC
- [SPARK-2750](#) - HTTPS support for History Server and web UI
- Added support for Spark SQL (DataFrames) in PySpark.
- Added support for the following MLlib features:
 - `spark.ml`
 - ML pipeline APIs

CDH 5 Release Notes

- The `hbase-spark` module has been added, which provides support for using HBase data in Spark jobs using `HBaseContext` and `JavaHBaseContext` contexts. See the [HBase and Spark chapter](#) of the Apache HBase Reference Guide for details about building a Spark application with HBase support.
- The `DATE` datatype is supported for Parquet and Avro tables.

Apache Sentry

[SENTRY-510](#) - Added support for collection of metrics for the Sentry/HDFS plugin

YARN

- Preemption guarantees that important tasks are not starved for resources, while allowing the CDH cluster to be used for experimental and research tasks. In FairScheduler, you can now disable preemption for a specific queue to ensure that its resources are not taken for other tasks.

What's New in CDH 5.7.1

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.7.1](#) on page 2277.

What's New in CDH 5.7.2

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.7.2](#) on page 2273.

What's New in CDH 5.7.3

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.7.3](#) on page 2272.

What's New in CDH 5.7.4

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.7.4](#) on page 2269.

What's New in CDH 5.7.5

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.7.5](#) on page 2267.

What's New in CDH 5.7.6

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.7.6](#) on page 2264.

What's New In CDH 5.6.x

What's New in CDH 5.6.0

CDH 5.6.0 is a maintenance release that introduces the EMC DSSD D5 Storage Appliance Integration for Hadoop DataNodes.

What's New in CDH 5.6.1

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.6.1](#) on page 2291.

What's New In CDH 5.5.x

What's New in CDH 5.5.0

The following sections describe new features introduced in CDH 5.5.0.

Operating System and Database Support

- **Operating Systems** - Support for RHEL/CentOS 6.6 (in SE Linux mode), 6.7, and 7.1, and Oracle Enterprise Linux 7.1.



Important: Cloudera supports RHEL 7 with the following limitations:

- Only RHEL 7.1 is supported. RHEL 7.0 is not supported.
- Only new installations of RHEL 7.1 are supported by Cloudera. For upgrades to RHEL 7.1, contact your OS vendor and see [Does Red Hat support upgrades between major versions of Red Hat Enterprise Linux?](#)

- **Databases** - Supports MariaDB 5.5, Oracle 12c, and PostgreSQL 9.4.

Apache Flume

- CDH 5.5 release is rebased on Flume 1.6.
- [FLUME-2498](#) Taildir source.
- [FLUME-2215](#) ResettableFileInputStream support for ucs-4 character.
- [FLUME-2729](#) PollableSource backoff times made configurable.
- [FLUME-2628](#) Netcat source support for different source encodings.
- [FLUME-2753](#) Support for empty replace string in Search and Replace interceptor.
- [FLUME-2763](#) Flume_env script support to handle JVM parameters.
- [FLUME-2095](#) JMS source support for username and password.

Apache Hadoop

- [HADOOP-1540](#) - DistCp supports file exclusions with a new filter option, `-exclusions <argument>`, to prevent files from being copied. The argument is a file that contains a list of Java regex patterns (one per line). If an exclusion pattern is matched, the file is not copied. To use, pass `-filters <pathToFileterFile>` to the `distcp` command.
- [HADOOP-8989](#) - The Hadoop shell now has a `find` utility, like that in UNIX, that allows users to search for files by name. Run `hadoop fs -help find` for more info.
- [HADOOP-11219](#), [HADOOP-7280](#) - WebImageViewer was upgraded to Netty 4. This does not affect the external classpath of Hadoop.
- [HADOOP-11827](#) - DistCp `buildListing()` now uses a threadpool to improve performance. To use, pass `--numListstatusThreads <numThreads>` to the `distcp` command. The default value is 1.
- [HDFS-6133](#) - HDFS balancer supports the exclusion of subtrees because running the HDFS balancer can destroy local data that is important for applications such as the HBase RegionServer.
- [HDFS-8828](#) - DistCp leverages HDFS snapshot diff to more easily build file and directory lists. The snapshot diff report provides diff information between two snapshots or between a snapshot and a non-HDFS directory.
- Improvements to HDFS scalability and performance:
 - [HDFS-7279](#) - In CDH 5.5.0 and higher, DataNode WebHDFS implementation uses Netty as an HTTP server instead of Jetty. With improved buffer and connection management, Netty lowers the risk for DataNode latency and OutOfMemoryError (OOM).
 - [HDFS-7435](#) and [HDFS-8867](#) add more efficient over-the-wire encoding.
 - [HDFS-7923](#) adds rate-limiting for block reports so that the NameNode is not swamped by DataNodes sending too many block reports at once.
 - [HDFS-7923](#) and [HDFS-7999](#) eliminate some cases on the DataNode side where I/O errors lead to scans being repeated on the local disks.
 - [HDFS-8581](#) fixes some cases where a lock is held for too long.
 - [HDFS-8792](#) and [HDFS-7609](#) optimize data structures on the NameNode side.
 - [HDFS-9107](#) fixes a bug that could limit scalability on larger clusters by causing the NameNode to falsely consider DataNodes to be dead.
 - Other bugs included: [HADOOP-11785](#), [HADOOP-12172](#), [HADOOP-11659](#), [HDFS-8845](#)

Apache HBase

- CDH now includes a *scanner heartbeat check*, which enforces a time limit on the execution of scan RPC requests. When the server receives a scan RPC request, a time limit is calculated to be half of the smaller of the two values `hbase.client.scanner.timeout.period` and `hbase.rpc.timeout`. When the time limit is reached, the server will return the results it has accumulated up to that point. For more information, see [Configuring the HBase Scanner Heartbeat](#).

Cloudera Search

- Cloudera Search adds support for Kerberos authentication for hosts running Solr behind a proxy server.
- Cloudera Search adds support for using LDAP and Active Directory for authentication.
- `solrctl` supports the `Config` API.

`solrctl` includes a `config` command that uses the `Config` API to directly manage configurations represented in `Config` objects. `Config` objects represent collection configuration information as specified by the `solrctl collection --create -c configName` command. `instancedir` and `Config` objects handle the same information, meeting the same need from the Solr server perspective, but there a number of differences between these two implementations.

Table 52: Config and instancedir Comparison

Attribute	Config	instancedir
Security	<p>Security support provided.</p> <ul style="list-style-type: none"> • In a Kerberos-enabled cluster, the ZooKeeper hosts associated with configurations created using the Config API automatically have proper ZooKeeper ACLs. • Sentry can be used to control access to the Config API, providing access control. For more information, see Configuring Sentry Authorization for Cloudera Search. 	<ul style="list-style-type: none"> • No ZooKeeper security support. Any user can create, delete, or modify an <code>instancedir</code> directly in ZooKeeper. • Because <code>instancedir</code> updates ZooKeeper directly, it is the client's responsibility to add the proper ACLs, which can be cumbersome.
Creation method	Generated from existing <code>config</code> or <code>instancedir</code> in ZooKeeper using the ConfigSet API.	Manually edited locally and re-uploaded directly to ZooKeeper using <code>solrctl</code> utility.
Template support	<ul style="list-style-type: none"> • Several predefined templates are available. These can be used as the basis for creating additional configs. Additional templates can be created by creating configs that are immutable. • Mutable templates that use a Managed Schema can be modified using the Schema API as opposed to being manually edited. As a result, configs are less flexible, but they are also less error-prone than <code>instancedirs</code>. 	One standard template.

Attribute	Config	instancedir
Sentry support	Configs include a number of templates, each with Sentry-enabled and non-Sentry-enabled versions. To enable Sentry, choose a Sentry-enabled template.	instancedirs include a single template that supports enabling Sentry. To enable Sentry with instancedirs, overwrite the original <code>solrconfig.xml</code> file with <code>solrconfig.xml.secure</code> as described in Enabling Solr as a Client for the Sentry Service Using the Command Line .

- Solr includes a set of built-in immutable configurations.

These templates are instantiated when Solr is initialized. This means these templates are not automatically available after an upgrade. To enable these templates on upgraded installations, use `solrctl init` or initialize Solr using Cloudera Manager. The newly included templates and the functionality each template supports are as follows:

Table 53: Available Config Templates and Attributes

Template Name	Supports Schema API	Uses Schemaless Solr	Supports Sentry
predefinedTemplate			
managedTemplate	■		
schemalessTemplate	■	■	
predefinedTemplateSecure			■
managedTemplateSecure	■		■
schemalessTemplateSecure	■	■	■

Apache Sentry (incubating)

- Sentry is rebased on Apache Sentry 1.5.1.
- Sentry introduces column-level access control for tables in Hive and Impala. Previously, Sentry supported privilege granularity only at the table level. To restrict access to a column of sensitive data, you needed to first create a view for a subset of columns, and then grant privileges on that view. Instead, Sentry now allows you to assign the `SELECT` privilege on a subset of columns in a table.
- Support for enabling Kerberos authentication for the Sentry web server.

Apache Spark

- Spark is rebased on Apache Spark 1.5.0.
- Dynamic allocation is enabled by default. You can explicitly disable dynamic allocation by using the option: `spark.dynamicAllocation.enabled = false`. Dynamic allocation is implicitly disabled if `--num-executors` is specified in the job.
- The following Spark libraries are now supported:
 - Spark SQL (including DataFrames). The following Spark SQL features are not supported:
 - Thrift JDBC/ODBC server
 - Spark SQL CLI
 - MLlib. The following MLlib features are not supported:
 - `spark.ml`
 - ML pipeline APIs
 -

CDH 5 Release Notes

Apache Sqoop

- Sqoop is rebased on Apache Sqoop 1.4.6.

What's New in CDH 5.5.1

This is a maintenance release that fixes important issues in Apache Commons and Apache HBase; for details, see [Issues Fixed in CDH 5.5.1](#) on page 2315.

What's New in CDH 5.5.2

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.5.2](#) on page 2310.

What's New in CDH 5.5.4

**Note:**

There is no CDH 5.5.3 release. This skip in the CDH 5.x sequence allows the CDH and Cloudera Manager components of Cloudera Enterprise 5.5.4 to have consistent numbering.

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.5.4](#) on page 2304.

What's New in CDH 5.5.5

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.5.5](#) on page 2301.

What's New in CDH 5.5.6

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.5.6](#) on page 2297.

What's New In CDH 5.4.x

What's New in CDH 5.4.0

**Important:**

Upgrading to CDH 5.4.0 and later from any earlier release requires an HDFS metadata upgrade. Be careful to follow all of the upgrade steps as instructed.

For the latest Impala features, see [New Features in Impala 2.2.x / CDH 5.4.x](#) on page 1886.

Operating System Support

CDH 5.4.0 adds support for RHEL and CentOS 6.6.

Security

The following summarizes new security capabilities in CDH 5.4.0:

- Secure Hue impersonation support for the Hue HBase application.
- Redaction of sensitive data from logs, centrally managed by Cloudera Manager, which prevents the `WHERE` clause in queries from leaking sensitive data into logs and management UIs.
- Cloudera Manager support for custom Kerberos principals.
- Kerberos support for Sqoop 2.
- Kerberos and TLS/SSL support for Flume Thrift source and sink.
- Navigator SAML support (requires Cloudera Manager).
- Navigator Key Trustee can now be installed and monitored by Cloudera Manager.
- Search can be configured to use SSL.

- Search supports protecting Solr and Lily HBase Indexer metadata using ZooKeeper ACLs in a Kerberos-enabled environment.

Apache Crunch

New HBase-related features:

- `HBaseTypes.cells()` was added to support serializing HBase Cell objects.
- All of the `HFileUtils` methods now support `PCollectionC extends Cell`, which includes both `PCollectionKeyValue` and `PCollectionCell`, on their method signatures.
- `HFileTarget`, `HBaseTarget`, and `HBaseSourceTarget` all support any subclass of `Cell` as an output type. `HFileSource` and `HBaseSourceTarget` still return `KeyValue` as the input type for backward compatibility with existing Crunch pipelines.

Developers can use Cell-based APIs in the same way as KeyValue-based APIs if they are not ready to update their code, but will probably have to change code inside `DoFns` because HBase 0.99 and later APIs deprecated or removed a number of methods from the HBase 0.96 API.

Apache Flume

CDH 5.4.0 adds SSL and Kerberos support for the Thrift source and sink, and implements DatasetSink 2.0.

Apache Hadoop

HDFS

- CDH 5.4.0 implements HDFS 2.6.0.
- CDH 5.4.0 HDFS provides hot-swap capability for DataNode disk drives. You can add or replace HDFS data volumes without shutting down the DataNode host ([HDFS-1362](#)).
- CDH 5.4.0 introduces cluster-wide redaction of sensitive data in logs and SQL queries. See [Sensitive Data Redaction](#).
- CDH 5.4.0 adds support for [Heterogenous Storage Policies](#).
- HDFS 2.6.0+ supports the option to configure AES encryption for block data transfer, using the property `dfs.encrypt.data.transfer.algorithm`. AES offers improved cryptographic strength and performance over the prior options of 3DES and RC4.

MapReduce

CDH 5.4.0 implements [MAPREDUCE-5785](#), which simplifies MapReduce job configuration. Instead of having to set both the heap size (`mapreduce.map.java.opts` or `mapreduce.reduce.java.opts`) and the container size (`mapreduce.map.memory.mb` or `mapreduce.reduce.memory.mb`), you can now choose to set only one of them; the other is inferred from `mapreduce.job.heap.memory-mb.ratio`. If you do not specify either of them, the container size defaults to 1 GB and the heap size is inferred.

For jobs that do not set the heap size, the JVM size increases from 200 MB to a default 820 MB. This is adequate for most jobs, but streaming tasks might need more memory because the Java process causes total usage to exceed the container size. This typically occurs only for those tasks relying on aggressive garbage collection to keep the heap under 200 MB.

YARN

- [YARN-2990](#) improves application launch time by 6 seconds when using FairScheduler (with the default Cloudera Manager settings shown in [YARN \(MR2 Included\) Properties in CDH 5.4.0](#)).

Apache HBase

CDH 5.4.0 implements HBase 1.0.

MultiWAL Support for HBase

CDH 5.4.0 introduces MultiWAL support for HBase region servers, allowing you to increase throughput when a region writes the write-ahead log (WAL).

doAs Impersonation for HBase

CDH 5.4.0 introduces `doAs` impersonation for the HBase Thrift server. `doAs` impersonation allows a client to authenticate to HBase as any user, and re-authenticate at any time, instead of as a static user only.

Read Replicas for HBase

CDH 5.4.0 introduces read replicas, along with a new timeline consistency model. This feature allows you to balance consistency and availability on a per-read basis, and provides a measure of high availability for reads if a RegionServer becomes unavailable.

Storing Medium Objects (MOBs) in HBase

CDH 5.4.0 HBase MOB allows you to store objects up to 10 MB (medium objects, or MOBs) directly in HBase while maintaining read and write performance.

Apache Hive

CDH 5.4.0 implements Hive 1.1.0. New capabilities include:

- A test-only version of [Hive on Spark](#) with the following limitations:
 - Parquet does not currently support vectorization; it simply ignores the setting of `hive.vectorized.execution.enabled`.
 - Hive on Spark does not yet support dynamic partition pruning.
 - Hive on Spark does not yet support HBase. If you want to interact with HBase, Cloudera recommends that you use Hive on MapReduce.

To deploy and test Hive on Spark in a test environment, use Cloudera Manager.



Important: Hive on Spark is included in CDH 5.4 and higher but is not currently supported nor recommended for production use. To try this feature, use it in a test environment until Cloudera resolves currently existing issues and limitations to make it ready for production use.

- Support for JAR files changes without scheduled maintenance.

To implement this capability, proceed as follows:

1. Set `hive.reloadable.aux.jars.path` in `/etc/hive/conf/hive-site.xml` to the directory that contains the JAR files.
2. Execute the `reload;` statement on HiveServer2 clients such as Beeline and the Hive JDBC.

- Beeline support for retrieving and printing query logs.

Some features in the upstream release are not yet supported for production use in CDH; these include:

- [HIVE-7935](#) - Support dynamic service discovery for HiveServer2
- [HIVE-6455](#) - Scalable dynamic partitioning and bucketing optimization
- [HIVE-5317](#) - Implement insert, update, and delete in Hive with full ACID support
- [HIVE-7068](#) - Integrate `AccumuloStorageHandler`
- [HIVE-7090](#) - Support session-level temporary tables in Hive
- [HIVE-7341](#) - Support for Table replication across HCatalog instances
- [HIVE-4752](#) - Add support for HiveServer2 to use Thrift over HTTP

Hue

CDH 5.4.0 adds the following:

- New Oozie editor
- Performance improvements
- New Search facets
- HBase impersonation

Kite

Kite in CDH has been rebased on the 1.0 release upstream. This breaks backward compatibility with existing APIs. The APIs are documented at <http://kitesdk.org/docs/1.0.0/apidocs/index.html>.

Notable changes are:

- Dataset writers that implement flush and sync now extend interfaces (Flushable and Syncable). Writers that no longer have misleading flush and sync methods.
- `DatasetReaderException`, `DatasetWriterException`, and `DatasetRepositoryException` have been removed and replaced with more specific exceptions, such as `IncompatibleSchemaException`. Exception classes now indicate what went wrong instead of what threw the exception.
- The `partition` API is no longer exposed; use the `view` API instead.
- `kite-data-hcatalog` is now `kite-data-hive`.



Note:

From 1.0 on, Kite will be strict about breaking compatibility and will use [semantic versioning](#) to signal which compatibility guarantees you can expect from a release (for example, incompatible changes require increasing the major version number). For more information, see the [Hello, Kite SDK 1.0](#) blog post.

Apache Oozie

- Added Spark action which lets you run Spark applications from Oozie workflows. See the [Oozie documentation](#) for more details.
- The Hive2 action now collects and reports Hadoop Job IDs for MapReduce jobs launched by Hive Server 2.
- The launcher job now uses YARN uber mode for all but the Shell action; this reduces the overhead (time and resources) of running these Oozie actions.

Apache Parquet

- Parquet memory manager now changes the row group size if the current size is expected to cause out-of-memory (OOM) errors because too many files are open. This causes a `WARN` message to be printed in the logs. A new setting, `parquet.memory.pool.ratio`, controls the percentage of the JVM's heap memory Parquet attempts to use.
- To improve job startup time, footers are no longer read by default for MapReduce jobs ([PARQUET-139](#)).



Note:

To revert to the old behavior (`ParquetFileReader` reads in all the files to obtain the footers), set `parquet.task.side.metadata` to `false` in the job configuration.

- The Parquet Avro object model can now read lists and maps written by Hive, Avro, and Thrift (similar capabilities were added to Hive in CDH 5.3). This compatibility fix does not change behavior. The extra record layer wrapping the list elements when Avro reads lists written by Hive can now be removed; to do this, set the expected Avro schema or set `parquet.avro.add-list-element-records` to `false`.
- Avro's map representation now writes null values correctly.
- The Parquet Thrift object model can now read data written by other object models (such as Hive, Impala, or Parquet-Avro), given a Thrift class for the data; compile a Thrift definition into an object, and supply it when creating the job.

Cloudera Search

- Solr metadata stored in ZooKeeper can now be protected by Zookeeper ACLs. In a Kerberos-enabled environment, Solr metadata stored in ZooKeeper is owned by the `solr` user and cannot be modified by other users.



Note:

- The Solr principal name can be configured in Cloudera Manager. The default name is `solr`, although other names can be specified.
- Collection configuration information stored under the `/solr/configs` znode is not affected by this change. As a result, collection configuration behavior is unchanged.

Administrators who modify Solr ZooKeeper metadata through operations like `solrctl init` or `solrctl cluster --put-solrxml` must now supply `solrctl` with a JAAS configuration using the `--jaas` configuration parameter. The JAAS configuration must specify the principal, typically `solr`, that the solr process uses.

End users, who typically do not need to modify Solr metadata, are unaffected by this change.

- Lily HBase Indexer metadata stored in ZooKeeper can now be protected by ZooKeeper ACLs. In a Kerberos-enabled environment, Lily HBase Indexer metadata stored in ZooKeeper is owned by the Solr user and cannot be modified by other users.

End users, who typically do not manage the Lily HBase Indexer, are unaffected by this change.

- The Lily HBase Indexer supports restricting access using Sentry.
- Services included with Search for CDH 5.4.0, including Solr, Key-Value Store Indexer, and Flume, now support SSL.
- The Spark Indexer and the Lily HBase Batch Indexer support delegation tokens for mapper-only jobs.
- Search for CDH 5.4.0 implements [SOLR-5746](#), which improves `solr.xml` file parsing. Error checking for duplicated options or unknown option names was added. These checks can help identify mistakes made during manual edits of the `solr.xml` file. User-modified `solr.xml` files may cause errors on startup due to these parsing improvements.
- By default, `CloudSolrServer` now uses multiple threads to add documents.



Note: Note: Due to multithreading, if document addition is interrupted by an exception, some documents, in addition to the one being added when the failure occurred, may be added.

To get the old, single-threaded behavior, set `parallel updates` to `false` on the `CloudSolrServer` instance.

Related JIRA: [SOLR-4816](#).

- Updates are routed directly to the correct shard leader, eliminating document routing at the server. This allows for near linear indexing throughput scalability. Document routing requires that the `solrj` client must know each document's unique identifier. The unique identifiers allow the client to route the update directly to the correct shard. For additional information, see [Shards and Indexing Data in SolrCloud](#).

Related JIRA: [SOLR-4816](#).

- The `loadSolr morphline` command supports nested documents. For more information, see [Morphlines Reference Guide](#).
- Navigator can be used to audit Cloudera Search activity.
- Search for CDH 5.4 supports logging queries before they are executed. This allows you can identify queries that could increase resource consumption. This also enables improving schemas or filters to meet your performance requirements. To enable this feature, set the `SolrCore` and `SolrCore.Request log` level to `DEBUG`.

Related JIRA: [SOLR-6919](#)

- `UniqFieldsUpdateProcessorFactory`, which Solr Server implements, has been improved to support all of the `FieldMutatingUpdateProcessorFactory` selector options. The `<lst named="fields"> init` param option is deprecated. Replace this option with `<arr name="fieldName">`.

If the `<lst named="fields"> init` param option is used, Solr logs a warning.

Related JIRA: [SOLR-4249](#).

- Configuration information was previously available using `FieldMutatingUpdateProcessorFactory` (`oneOrMany` or `getBooleanArg`). Those methods are now deprecated. The methods have been moved to `NamedList` and renamed to `removeConfigArgs` and `removeBooleanArg`, respectively.

If the `oneOrMany` or `getBooleanArg` methods of `FieldMutatingUpdateProcessorFactory` are used, Solr logs a warning.

Related JIRA: [SOLR-5264](#).

Apache Spark

CDH 5.4.0 Spark is rebased on Apache Spark 1.3.0 and provides the following new capabilities:

- Spark Streaming WAL (write-ahead log) on HDFS, preventing any data loss on driver failure
- Kafka connector for Spark Streaming to avoid the need for the HDFS WAL
- Spark Streaming recovery is supported for production use
- Spark external shuffle service
- Improvements in automatically setting CDH classpaths for Avro, Parquet, Flume, and Hive
- Improvements in the collection of task metrics

The following is not supported in a production environment because of its immaturity:

- Spark SQL (which now includes dataframes)

See also [Apache Spark Known Issues](#) on page 2043 and [Apache Spark Incompatible Changes and Limitations](#) on page 1942.

Apache Sqoop

- **Sqoop 2:**
 - Implements Sqoop 2 version 1.99.5.
 - Sqoop 2 supports Kerberos.
 - Sqoop 2 supports PostgreSQL as the repository database.

What's New in CDH 5.4.1

This is a maintenance release that fixes the following issue; for details of other important fixes, see [Issues Fixed in CDH 5.4.1](#) on page 2339.

Upgrades to CDH 5.4.1 from Releases Earlier than 5.4.0 May Fail

Problem: Because of a change in the implementation of the NameNode metadata upgrade mechanism, upgrading to CDH 5.4.1 from a version lower than 5.4.0 can take an inordinately long time. In a cluster with NameNode high availability (HA) configured and a large number of edit logs, the upgrade can fail, with errors indicating a timeout in the pre-upgrade step on JournalNodes.

What to do:

To avoid the problem: Do not upgrade to CDH 5.4.1; upgrade to CDH 5.4.2 instead.

If you experience the problem: If you have already started an upgrade and seen it fail, contact Cloudera Support. This problem involves no risk of data loss, and manual recovery is possible.

If you have already completed an upgrade to CDH 5.4.1, or are installing a new cluster: In this case you are not affected and can continue to run CDH 5.4.1.

Cloudera Search

- Beginning with CDH 5.4.1, Search for CDH supports configurable transaction log replication levels for replication logs stored in HDFS.

Apache Spark

- Spark supports submitting python applications in cluster mode.

What's New in CDH 5.4.2

This is a maintenance release that fixes the following issue; for details of other important fixes, see [Issues Fixed in CDH 5.4.2](#) on page 2339.

Upgrades to CDH 5.4.1 from Releases Earlier than 5.4.0 May Fail

Problem: Because of a change in the implementation of the NameNode metadata upgrade mechanism, upgrading to CDH 5.4.1 from a version lower than 5.4.0 can take an inordinately long time. In a cluster with NameNode high availability (HA) configured and a large number of edit logs, the upgrade can fail, with errors indicating a timeout in the pre-upgrade step on JournalNodes.

What to do:

CDH 5 Release Notes

To avoid the problem: Do not upgrade to CDH 5.4.1; upgrade to CDH 5.4.2 instead.

If you experience the problem: If you have already started an upgrade and seen it fail, contact Cloudera Support. This problem involves no risk of data loss, and manual recovery is possible.

If you have already completed an upgrade to CDH 5.4.1, or are installing a new cluster: In this case you are not affected and can continue to run CDH 5.4.1.

What's New in CDH 5.4.3

This is a maintenance release that fixes the following issue; for details of other important fixes, see [Issues Fixed in CDH 5.4.3](#) on page 2337.

NameNode Incorrectly Reports Missing Blocks During Rolling Upgrade

Problem: During a rolling upgrade to any of the releases listed below, the NameNode may report missing blocks after rolling back multiple DataNodes. This is caused by a race condition with block reporting between the DataNode and the NameNode. No permanent data loss occurs, but data can be unavailable for up to six hours before the problem corrects itself.

Releases affected: CDH 5.0.6, 5.1.5, 5.2.5, 5.3.3, 5.4.1, 5.4.2

What to do:

To avoid the problem: Cloudera advises skipping the affected releases and installing a release containing the fix. For example, do not upgrade to CDH 5.4.2; upgrade to CDH 5.4.3 instead.

The releases containing the fix are: CDH 5.3.4, 5.4.3

If you have already completed an upgrade to an affected release, or are installing a new cluster: You can continue to run the release, or upgrade to a release that is not affected.

What's New in CDH 5.4.4

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.4.4](#) on page 2336.

What's New in CDH 5.4.5

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.4.5](#) on page 2333.

What's New in CDH 5.4.7



Note:

There is no CDH 5.4.6 release. This skip in the CDH 5.x sequence allows the CDH and Cloudera Manager components of Cloudera Enterprise 5.4.7 to have consistent numbering.

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.4.7](#) on page 2331

What's New in CDH 5.4.8

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.4.8](#) on page 2328

What's New in CDH 5.4.9

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.4.9](#) on page 2326.

What's New in CDH 5.4.10

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.4.10](#) on page 2323.

What's New in CDH 5.4.11

This is a maintenance release that fixes some important issues. For details, see [Issues Fixed in CDH 5.4.11](#) on page 2320

What's New In CDH 5.3.x

What's New in CDH 5.3.0

Oracle JDK 8 Support

CDH 5.3 supports Oracle JDK 1.8.

Apache Hadoop

HDFS

CDH 5.3 provides the following new capabilities:

- **HDFS Data At Rest Encryption** - This feature is now ready for use in production environments.



Important:

Client hosts may need a more recent version of `libcrypto.so`. See [Apache Hadoop Known Issues](#) on page 1952 for more information.



Important: Cloudera provides two solutions:

- **Navigator Encrypt** is production ready and available to Cloudera customers licensed for Cloudera Navigator. Navigator Encrypt operates at the Linux volume level, so it can encrypt cluster data inside and outside HDFS. Consult your Cloudera account team for more information.
- **HDFS Encryption** is production ready and operates at the HDFS directory level, enabling encryption to be applied only to HDFS folders where needed.

- **S3A** - S3A is an HDFS implementation of the Simple Storage Service (S3) from Amazon Web Services. It is similar to S3N, which is the other implementation of this functionality. The key difference is that S3A relies on the officially-supported AWS Java SDK for communicating with S3, while S3N uses a best-effort-supported `jets3t` library to do the same. For a listing of the parameters, see [HADOOP-10400](#).

YARN

YARN now provides a way for long-running applications to get new delegation tokens.

Apache Flume

CDH 5.3 provides a Kafka Channel ([FLUME-2500](#)).

Apache HBase

CDH 5.3 provides `checkAndMutate(RowMutations)`, in addition to existing support for atomic `checkAndPut` as well as `checkAndDelete` operations on individual rows ([HBASE-11796](#)).

Apache Hive

- Hive can use multiple HDFS encryption zones.
- Hive-HBase integration contains many fixes and new features such as reading HBase snapshots.
- Many Hive Parquet fixes.
- Hive Server 2 can handle multiple LDAP domains for authentication.

Hue

New Features:

- Hue is re-based on Hue 3.7
- SAML authentication has been revamped.
- CDH 5.3 simplifies the task of configuring Hue to store data in an Oracle database by bundling the Oracle Install Client.

Apache Oozie

- You can now update the definition and properties of an already running Coordinator. See the [documentation](#) for more information.
- A new `poll` command in the Oozie client polls a Workflow Job, Coordinator Job, Coordinator Action, or Bundle Job until it finishes. See the [documentation](#) for more information.

Apache Parquet

- [PARQUET-132](#): Add type parameter to AvroParquetInputFormat for Spark
- [PARQUET-107](#): Add option to disable summary metadata files
- [PARQUET-64](#): Add support for new type annotations (date, time, timestamp, etc.)

Cloudera Search

New Features:

- Cloudera Search includes a version of Kite 0.15.0, which includes all morphlines-related backports of all fixes and features in Kite 0.17.1. Morphlines now includes functionality that enables partially updating document as well as deleting documents. Partial updating or deleting can be completed by unique IDs or by documents that match a query. For additional information on Kite, see:
 - [Kite repository](#)
 - [Kite Release Notes](#)
 - [Kite documentation](#)
 - [Kite examples](#)
- CrunchIndexerTool now sends a commit to Solr on job success.
- Added support for deleting documents stored in Solr [by unique id](#) as well as [by query](#).

Apache Sentry (incubating)

- Sentry HDFS Plugin - Allows you to configure synchronization of Sentry privileges to HDFS ACLs for specific HDFS directories. This simplifies the process of sharing table data between Hive or Impala and other clients (such as MapReduce, Pig, Spark), by automatically updating the ACLs when a `GRANT` or `REVOKE` statement is executed. It also allows all roles and privileges to be managed in a central location (by Sentry).
- Metrics - CDH 5.3 supports metrics for the Sentry service. These metrics can be reported either through JMX or the console; configure this by setting the property `sentry.service.reporter` to `jmx` or `console`. A Sentry web server listening by default on port 51000 can expose the metrics in `json` format. Web reporting is disabled by default; enable it by setting `sentry.service.web.enable` to `true`. You can configure the port on which Sentry web server listens by means of the `sentry.service.web.port` property .

Apache Spark

- CDH Spark has been rebased on Apache Spark 1.2.0.
- Spark Streaming can now save incoming data to a WAL (write-ahead log) on HDFS, preventing any data loss on driver failure.



Important:

This feature is currently in Beta; Cloudera includes it in CDH Spark but does not support it.

- The YARN back end now supports dynamic allocation of executors. See [Job Scheduling](#) for more information.
- Native library paths (set via Spark configuration options) are correctly propagated to executors in YARN mode ([SPARK-1719](#)).
- The Snappy codec should now work out-of-the-box on Linux distributions with older `glibc` versions such as CentOS 5.
- Spark SQL now includes the Spark Thrift Server in CDH.

**Important:**

Spark SQL remains an experimental and unsupported feature in CDH.

See [Apache Spark Incompatible Changes and Limitations](#) on page 1942 and [Apache Spark Known Issues](#) on page 2043 for additional important information.

Apache Sqoop

- **Sqoop 1:**

- The MySQL connector now fetches on a row-by row-basis.
- The SQL server now has **upsert** (insert or update) support ([SQOOP-1403](#)).
- The Oracle direct connector now works with index-organized tables ([SQOOP-1632](#)). To use this capability, you must set the chunk method to `PARTITION`:

```
-Doraoop.chunk.method=PARTITION
```

- **Sqoop 2:**

- FROM/TO re-factoring is now supported ([SQOOP-1367](#)).

What's New in CDH 5.3.1

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.3.1](#) on page 2360.

What's New in CDH 5.3.2

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.3.2](#) on page 2358.

What's New in CDH 5.3.3

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.3.3](#) on page 2356.

What's New in CDH 5.3.4

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.3.4](#) on page 2354.

What's New in CDH 5.3.5

This is a maintenance release that fixes the following issue. For details of other fixes, see [Issues Fixed in CDH 5.3.5](#) on page 2354:

Potential job failures during YARN rolling upgrades to CDH 5.3.4

Problem: A MapReduce security fix introduced a compatibility issue that results in job failures during YARN rolling upgrades from CDH 5.3.3 to CDH 5.3.4.

Cloudera Bug: CDH-28680

Release affected: CDH 5.3.4

Release containing the fix: CDH 5.3.5

Workarounds: You can use any one of the following workarounds for this issue:

- Upgrade to CDH 5.3.5.
- Restart any jobs that might have failed during the upgrade.
- Explicitly set the version of MapReduce to be used so it is picked on a per-job basis.
 1. Update the YARN property, **MR Application Classpath** (`mapreduce.application.classpath`), either in Cloudera Manager or in the `mapred-site.xml` file. Remove all existing values and add a new entry: `<parcel-path>/lib/hadoop-mapreduce/*`, where `<parcel-path>` is the absolute path to the parcel

installation. For example, the default installation path for the CDH 5.3.3 parcel would be:
`/opt/cloudera/parcels/CDH-5.3.3-1.cdh5.3.3.p0.5/lib/hadoop-mapreduce/*`.

2. Wait until jobs submitted with the above client configuration change have run to completion.
3. Upgrade to CDH 5.3.4.
4. Update the **MR Application Classpath** (`mapreduce.application.classpath`) property to point to the new CDH 5.3.4 parcel.

Do *not* delete the old parcel until after all jobs submitted prior to the upgrade have finished running.

What's New in CDH 5.3.6

This is a maintenance release that fixes several issues. For details, see [Issues Fixed in CDH 5.3.6](#) on page 2352.

What's New in CDH 5.3.8

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.3.8](#) on page 2350.

What's New in CDH 5.3.9

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.3.9](#) on page 2348.

What's New in CDH 5.3.10

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.3.10](#) on page 2346.

What's New In CDH 5.2.x

What's New in CDH 5.2.0

**Important:**

Upgrading to CDH 5.2.0 and later from any earlier release requires an HDFS metadata upgrade and other steps not usually required for a minor-release upgrade.

CDH 5.2.0 is a minor release which includes new features and bug fixes.

Go to [Issues Fixed In CDH 5.2.0](#) or keep reading for [New Features and Changes](#) on page 1848.

See also [Known Issues in CDH 5](#) on page 1944.

New Features and Changes

New features and changes are grouped by area.

Operating System Support

CDH 5.2.0 adds support for Ubuntu Trusty (version 14.04).

**Important:**

Installing CDH by adding a repository entails an additional step on Ubuntu Trusty, to ensure that you get the CDH version of ZooKeeper, rather than the version that is bundled with Trusty.

Apache Avro

CDH 5.2 implements Avro version 1.7.6, with backports from 1.7.7. Important changes include:

- [AVRO-1398](#): Increase default sync interval from 16k to 64k. There is a very small chance this could causes an incompatibility in some cases, but you can be control the interval by setting `avro.mapred.sync.interval` in the MapReduce job configuration. For example, set it to 16000 to get the old behavior.

- [AVRO-1355](#): Record schema should reject duplicate field names. This change rejects schemas with duplicate field names. This could affect some applications, but if schemas have duplicate field names then they are unlikely to work properly in any case. The workaround is to make sure a record's field names are unique within the record.

Apache Hadoop

HDFS

CDH 5.2 provides the following new capabilities:

- **HDFS Data at Rest Encryption**



Note: Cloudera provides the following two solutions for data at rest encryption:

- **Navigator Encrypt** - Is production ready and available for Cloudera customers licensed for Cloudera Navigator. Navigator Encrypt operates at the Linux volume level, so it can encrypt cluster data inside and outside HDFS. Talk to your Cloudera account team for more information about this capability.
- **HDFS Encryption** - Included in CDH 5.2.0 and operates at the HDFS folder level, enabling encryption to be applied only to HDFS folders where needed. This feature has several known limitations. Therefore, Cloudera does not currently support this feature in CDH 5.2 and it is *not* recommended for production use. To try the feature, upgrade to the latest version of CDH 5.

HDFS now implements transparent, end-to-end encryption of data read from and written to HDFS by creating encryption zones. An encryption zone is a directory in HDFS with all of its contents, that is, every file and subdirectory in it, encrypted. You can use either the **KMS** or the **Key Trustee** service to store, manage, and access encryption zone keys.

HDFS now implements transparent, end-to-end encryption of data read from and written to HDFS by creating encryption zones. An encryption zone is a directory in HDFS with all of its contents, that is, every file and subdirectory in it, encrypted.

- Extended attributes: HDFS `XATTRS` allow extended attributes to be stored per file (<https://issues.apache.org/jira/browse/HDFS-2006>).
- Authentication improvements when using an HTTP proxy server.
- A new Hadoop Metrics sink that allows writing directly to Graphite.
- Specification for Hadoop Compatible Filesystem effort.
- `OfflineImageViewer` to browse an `fsimage` via the WebHDFS API.
- Supportability improvements and bug fixes to the NFS gateway.
- Modernized web UIs (HTML5 and JavaScript) for HDFS daemons.

MapReduce

CDH 5.2 provides an optimized implementation of the mapper side of the MapReduce shuffle. The optimized implementation may require tuning different from the original implementation, and so it is considered experimental and is not enabled by default.

You can select this new implementation on a per-job basis by setting the job configuration value

```
mapreduce.job.map.output.collector.class to
org.apache.hadoop.mapred.nativetask.NativeMapOutputCollectorDelegator, or use enable Cloudera
Manager to enable it.
```

Some jobs which use custom writable types or comparators may not be able to take advantage of the optimized implementation.

the following new capabilities and improvements:

YARN

CDH 5.2 provides the following new capabilities and improvements:

- New features and improvements in the Fair Scheduler:
 - New features:
 - Fair Scheduler now allows setting the `fairsharePreemptionThreshold` per queue (leaf and non-leaf). This threshold is a decimal value between 0 and 1; if a queue's usage is under (`preemption-threshold * fairshare`) for a configured duration, resources from other queues are preempted to satisfy this queue's request. Set this value in `fair-scheduler.xml`. The default value is 0.5.
 - Fair Scheduler now allows setting the `fairsharePreemptionTimeout` per queue (leaf and non-leaf). For a starved queue, this timeout determines when to trigger preemption from other queues. Set this value in `fair-scheduler.xml`.
 - Fair Scheduler now shows the **Steady Fair Share** in the Web UI. The Steady Fair Share is the share of the cluster resources a particular queue or pool would get if all existing queues had running applications.
 - Improvements:
 - Fair Scheduler uses **Instantaneous Fair Share** (`fairshare` that considers only active queues) for scheduling decisions to improve the time to achieve steady state (fairshare).
 - The default for `maxAMShare` is now 0.5, meaning that only half the cluster's resources can be taken up by Application Masters. You can change this value in `fair-scheduler.xml`.
- YARN's REST APIs support submitting and killing applications.
- YARN's timeline store is integrated with Kerberos.

Apache Crunch

CDH 5.2 provides the following new capabilities:

- Improvements in [Scrunch](#), including:
 - New `join` API that matches the one in Crunch
 - New aggregation API, including support for [Algebird](#)-based aggregations
 - Built-in serialization support for all tuple types as well as case classes.
- A new module, `crunch-hive`, for reading and writing **Optimized Row Columnar** (ORC) Files with Crunch.

Apache Flume

CDH 5.2 provides the following new capabilities:

- [Kafka](#) Integration: Flume can now accept data from Kafka via the `KafkaSource` ([FLUME-2250](#)) and push to Kafka using the `KafkaSink` ([FLUME-2251](#)).
- Kite Sink can now write to Hive and HBase datasets ([FLUME-2463](#)).
- Flume agents can now be configured via Zookeeper (experimental, [FLUME-1491](#))
- Embedded Agents now support Interceptors ([FLUME-2426](#))
- `syslog` Sources now support configuring which fields should be kept ([FLUME-2438](#))
- File Channel replay is now much faster ([FLUME-2450](#))
- New regular-expression search-and-replace interceptor ([FLUME-2431](#))
- Backup checkpoints can be optionally compressed ([FLUME-2401](#))

Hue

CDH 5.2 provides the following new capabilities:

- New application for editing Sentry roles and Privileges on databases and tables
- Search App
- Heatmap, Tree, Leaflet widgets
- Micro-analysis of fields
- Exclusion facets
- Oozie Dashboard: bulk actions, faster display
- File Browser: drag-and-drop upload, history, ACLs edition

- Hive and Impala: LDAP pass-through, query expiration, TLS/SSL (Hive), new graphs
- Job Browser: YARN kill application button

Apache HBase

CDH 5.2 implements HBase 0.98.6, which represents a minor upgrade to HBase. This upgrade introduces new features and moves some features which were previously marked as experimental to fully supported status.

Apache Hive

CDH 5.2 introduces the following important changes in Hive.

- CDH 5.2 implements Hive 0.13, providing the following new capabilities:
 - Sub-queries in the `WHERE` clause
 - Common table expressions (CTE)
 - Parquet supports `timestamp`
 - HiveServer2 can be configured with a `hiverc` file that is automatically run when users connect
 - Permanent UDFs
 - HiveServer2 session and operation timeouts
 - Beeline accepts a `-i` option to initialize with a SQL file
 - New `join` syntax (implicit joins)
- As of CDH 5.2.0, you can create Avro-backed tables simply by using `STORED AS AVRO` in a DDL statement. The AvroSerDe takes care of creating the appropriate Avro schema from the Hive table schema, making it much easier to use Avro with Hive.
- Hive supports additional datatypes, as follows:
 - Hive can read `char` and `varchar` datatypes written by Hive, and `char` and `varchar` datatypes written by Impala.
 - Impala can read `char` and `varchar` datatypes written by Hive and Impala.

These new types have been enabled by expanding the supported DDL, so they are backward compatible. You can add `varchar(n)` columns by creating new tables with that type, or changing a *string* column in existing tables to `varchar`.



Note:

`char(n)` columns are not stored in a fixed-length representation, and do not improve performance (as they do in some other databases). Cloudera recommends that in most cases you use `text` or `varchar` instead.

- `DESCRIBE DATABASE` returns additional fields: `owner_name` and `owner_type`. The command will continue to behave as expected if you identify the field you're interested in by its (string) name, but could produce unexpected results if you use a numeric index to identify the field(s).

Impala

Impala in CDH 5.2.0 includes major new features such as spill-to-disk for memory-intensive queries, subquery enhancements, analytic functions, and new `CHAR` and `VARCHAR` data types. For the full feature list and more details, see [What's New in Apache Impala](#) on page 1869.

Kite

Kite is an open source set of libraries, references, tutorials, and code samples for building data-oriented systems and applications. For more information about Kite, see the [Kite SDK Development Guide](#).

Kite has been rebased to version 0.15.0 in CDH 5.2.0, from the base version 0.10.0 in CDH 5.1. `kite-morphlines` modules are backward-compatible, but this change breaks backward-compatibility for the `kite-data` API.

Kite Data

The Kite data API has had substantial updates since the version included in CDH 5.1.

Changes from 0.15.0

The Kite version in CDH 5.2 is based on 0.15.0, but includes some newer changes. Specifically, it includes support for dataset namespaces, which can be used to set the database in the Hive Metastore.

The introduction of namespaces changed the file system repository layout; now there is an additional namespace directory for datasets stored in HDFS (`repository/namespace/dataset/`). There are no compatibility problems when you use `Dataset` URIs, but all datasets created with the `DatasetRepository` API will be located in a namespace directory. This new directory level is not expected in Kite 0.15.0 or 0.16.0 and will prevent the dataset from being loaded. The work-around is to switch to using `Dataset` URIs (see below) that include the namespace component. Existing datasets will work without modification.

Except as noted above, Kite 0.15.0 in CDH 5.2 is fully backward-compatible. It can load datasets written with any previous Kite version.

Dataset URIs

Datasets are identified with a single URI, rather than a repository URI and dataset name. The dataset URI contains all the information Kite needs to determine which implementation (Hive, HBase, or HDFS) to use for the dataset, and includes both the dataset's name and a namespace.

The Kite API has been updated so that developers call methods in the `Datasets` utility class as they would use `DatasetRepository` methods. The `Datasets` methods are recommended, and the `DatasetRepository` API is deprecated.

Views

The Kite data API now allows you to select a view of the dataset by setting constraints. These constraints are used by Kite to automatically prune unnecessary partitions and filter records.

MapReduce input and output formats

The `kite-data-mapreduce` module has been added. It provides both `DatasetKeyInputFormat` and `DatasetKeyOutputFormat` that allow you to run MapReduce jobs over datasets or views. Spark is also supported by the input and output formats.

Dataset CLI tool

Kite now includes a command-line utility that can run common maintenance tasks, like creating a dataset, migrating a dataset's schema, copying from one dataset to another, and importing CSV data. It also has helpers that can create Avro schemas from data files and other Kite-related configuration.

Flume DatasetSink

The Flume `DatasetSink` has been updated for the `kite-data` API changes. It supports all previous configurations without modification.

In addition, the `DatasetSink` now supports dataset URIs with the configuration option `kite.dataset.uri`.

Apache Mahout

Mahout jobs launched from the `bin/mahout` script will now use cluster's default parameters, rather than hard-coded parameters from the library. This may change the algorithms' run-time behavior, possibly for the better. ([MAHOUT-1565](#).)

Apache Oozie

CDH 5.2 introduces the following important changes:

- A new [Hive 2 Action](#) allows Oozie to run HiveServer2 scripts. Using the Hive Action with HiveServer2 is now deprecated; you should switch to the new Hive 2 Action as soon as possible.
- The MapReduce action can now also be configured by Java code
This gives users the flexibility of using their own driver Java code for configuring the MR job, while also getting the advantages of the MapReduce action (instead of using the Java action). See the [documentation](#) for more info.
- The `PurgeService` is now able to remove completed child jobs from long running coordinator jobs
- `ALL` can now be set for `oozie.service.LiteWorkflowStoreService.user.retry.error.code.ext` to make Oozie retry actions automatically for every type of error
- All Oozie servers in an Oozie HA group now synchronize on the same randomly generated rolling secret for signing auth tokens
- You can now upgrade from CDH 4.x to CDH 5.2 and later with jobs in `RUNNING` and `SUSPENDED` states. (An upgrade from CDH 4.x to a CDH 5.x release *earlier* than CDH 5.2.0 would still require that no jobs be in either of those states).

Apache Parquet (incubating)

CDH 5.2 Parquet is rebased on Parquet 1.5 and Parquet-format 2.1.0.

Cloudera Search

New Features:

- Cloudera Search adds support for Spark indexing using the `CrunchIndexerTool`.
- Cloudera Search adds fault tolerance for single-shard deployments. This fault tolerance is enabled with a new `-a` option in `solrctl`, which configures shards to automatically be re-added on an existing, healthy node if the node hosting the shard become unavailable.
- Components of Cloudera Search include Kite 0.15.0. This includes all morphlines-related backports of all fixes and features in Kite 0.17.0. For additional information on Kite, see:
 - [Kite repository](#)
 - [Kite Release Notes](#)
 - [Kite documentation](#)
 - [Kite examples](#)
- Search adds support for multi-threaded faceting on fields. This enables parallelizing operations, allowing them to run more quickly on highly concurrent hardware. This is especially helpful in cases where faceting operations apply to large datasets over many fields.
- Search adds support for distributed pivot faceting, enabling faceting on multi-shard collections.

Apache Sentry (incubating)

CDH 5.2 introduces the following changes to Sentry.

Sentry Service:

- If you are using the database-backed Sentry service, upgrading from CDH 5.1 to CDH 5.2 will require a schema upgrade.
- **Hive SQL Syntax:**
 - `GRANT` and `REVOKE` statements have been expanded to include `WITH GRANT OPTION`, thus allowing you to delegate granting and revoking privileges.
 - The `SHOW GRANT ROLE` command has been updated to allow non-admin users to list grants for roles that are currently assigned to them.
 - The `SHOW ROLE GRANT GROUP <groupName>` command has been updated to allow non-admin users that are part of the group specified by `<groupName>` to list all roles assigned to this group.

Apache Spark

CDH 5.2 Spark is rebased on Apache Spark/Streaming 1.1 and provides the following new capabilities:

- Stability and performance improvements.
- New sort-based shuffle implementation (disabled by default).
- Better performance monitoring through the Spark UI.
- Support for arbitrary Hadoop `InputFormats` in PySpark.
- Improved Yarn support with several bug fixes.

Apache Sqoop

CDH 5.2 Sqoop 1 is rebased on Sqoop 1.4.5 and includes the following changes:

- Mainframe connector added.
- Parquet support added.

There are no changes for Sqoop 2.

What's New in CDH 5.2.1

CDH 5.2.1 maintenance release that fixes the “POODLE” and Apache Hadoop Distributed Cache vulnerabilities described below. All CDH 5.2.0 users should upgrade to 5.2.1 as soon as possible.

Go to [Issues Fixed In CDH 5.2.1](#).

“POODLE” Vulnerability on TLS/SSL enabled ports

The POODLE (Padding Oracle On Downgraded Legacy Encryption) attack forces the use of the obsolete SSLv3 protocol and then exploits a cryptographic flaw in SSLv3. The only solution is to disable SSLv3 entirely. This requires changes across a wide variety of components of CDH and Cloudera Manager in 5.2.0 and all earlier versions. CDH 5.2.1 provides these changes for CDH 5.2.0 deployments. For more information, see the [Cloudera Security Bulletin](#).

Apache Hadoop Distributed Cache Vulnerability

The Distributed Cache Vulnerability allows a malicious cluster user to expose private files owned by the user running the YARN NodeManager process. For more information, see the [Cloudera Security Bulletin](#).

What's New in CDH 5.2.3

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.2.3](#).

What's New in CDH 5.2.4

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.2.4](#).

What's New in CDH 5.2.5

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.2.5](#).

What's New in CDH 5.2.6

This is a maintenance release that fixes some important issues; for details, see [Issues Fixed in CDH 5.2.6](#).

What's New in CDH 5.1.x

The following sections describe what's new in CDH 5.1.x releases. You can also review [Known Issues in CDH 5](#) on page 1944 and [Incompatible Changes and Limitations](#) on page 1916.

What's New in CDH 5.1.0

This is a minor release, which includes new features, changes, and fixed issues. See also [Issues Fixed in CDH 5.1.0](#) on page 2375.

New Features and Changes in CDH 5.1.0

This is a minor release which introduces the following new features and changes, organized by component. See also [What's New in Apache Impala](#) on page 1869.

Operating System Support

CDH 5.1 adds support for version 6.5 of RHEL and related platforms.

Apache Crunch

- CDH 5.1.0 implements Crunch 0.10.0.

Apache Flume

- CDH 5.1.0 implements Flume 1.5.0.

Apache Hadoop

HDFS

POSIX Access Control Lists: As of CDH 5.1, HDFS supports POSIX Access Control Lists (ACLs), an addition to the traditional POSIX permissions model already supported. ACLs provide fine-grained control of permissions for HDFS files by providing a way to set different permissions for specific named users or named groups.

NFS Gateway Improvements: CDH 5.1 makes the following improvements to the HDFS NFS gateway capability:

- Subdirectory mounts :
 - Previously, clients could mount only the HDFS root directory.
 - As of CDH 5.1, a single mount point, configured via the `nfs.export.point` property in `hdfs-site.xml` on the NFS gateway node, is available to clients.
- Improved support for Kerberized clusters ([HDFS-5898](#)):
 - Previously the NFS Gateway could connect to a secure cluster, but didn't support logging in from a keytab.
 - As of CDH 5.1, set the `nfs.kerberos.principal` and `nfs.keytab.file` properties in `hdfs-site.xml` to allow users to log in from a keytab.
- Support for port monitoring ([HDFS-6406](#)):
 - Previously, the NFS Gateway would always accept connections from any client.
 - As of CDH 5.1, set `nfs.port.monitoring.disabled` to `false` in `hdfs-site.xml` to allow connections only from privileged ports (those with root access).
- Static uid/gid mapping for NFS clients that are not in synch with the NFS Gateway ([HDFS-6435](#)):
 - NFS sends UIDs and GIDs over the network from client to server, meaning that the UIDs and GIDs must be in synch between clients and server machines in order for users and groups to be set appropriately for file access and file creation; this is usually but not always the case.
 - As of CDH 5.1, you can configure a static UID/GID mapping file, by default `/etc/nfs.map`.
 - You can change the default (to use a different file path) by means of the `nfs.static.mapping.file` property in `hdfs-site.xml`.
 - The following sample entries illustrate the format of the file:

```
uid 10 100 # Map the remote UID 10 the local UID 100
gid 11 101 # Map the remote GID 11 to the local GID 101
```

- Hadoop portmap, or insecure system portmap, no longer required:
 - Many supported OS have portmap bugs detailed [here](#).
 - CDH 5.1 allows you to circumvent the problems by starting the NFS gateway as root, whether you install CDH from packages or parcels.

**Note:**

After initially registering with the system portmap as root, the NFS Gateway drops privileges and runs as a regular user.

- Cloudera Manager starts the gateway as root by default.
- Support for AIX NFS clients ([HDFS-6549](#)):
 - To deploy AIX NFS clients, set `nfs.aix.compatibility.mode.enabled` to `true` in `hdfs-site.xml`.
 - This enables code that handles bugs in the AIX implementation of NFS.

MapReduce and YARN

YARN with Impala supports Dynamic Prioritization.

Apache HBase

- CDH 5.1.0 implements HBase 0.98.
- As of CDH 5.1.0, HBase fully supports BucketCache, which was introduced as an experimental feature in CDH 5 Beta 1.
- HBase now supports access control for `EXEC` permissions.
- CDH 5.1.0 HBase introduces a reverse scan API; allowing you to scan a table in reverse.
- You can now run a MapReduce job over a snapshot from HBase, rather than being limited to live data.
- A new stateless streaming scanner is available over the REST API.
- The `delete*` methods of the Delete class of the HBase Client API now use the timestamp from the constructor, the same behavior as the Put class. (In HBase versions before CDH 5.1, the `delete*` methods ignored the constructor's timestamp, and used the value of `HConstants.LATEST_TIMESTAMP`. This behavior was different from the behavior of the `add()` methods of the Put class.)
- The `SnapshotInfo` tool has been enhanced in the following ways:
 - A new option, `-list-snapshots`, has been added to the `SnapshotInfo` command. This option allows you to list snapshots on either a local or remote server.
 - You can now pass the `-size-in-bytes` flag to print the size of snapshot files in bytes rather than the default human-readable format.
 - The size of each snapshot file in bytes is checked against the size reported in the manifest, and if the two sizes differ, the tool reports the file as corrupt.
- A new `-target` option for `ExportSnapshot` allows you to specify a different name for the target cluster from the snapshot name on the source cluster.

In addition, Cloudera has fixed some binary incompatibilities between HBase 0.96 and 0.98. As a result, the incompatibilities introduced by [HBASE-10452](#) and [HBASE-10339](#) *do not* affect CDH 5.1 HBase, as explained below:

- HBASE-10452 introduced a new exception and error message in `setTimeStamp()`, for an extremely unlikely event when where getting a `TimeRange` could fail because of an integer overflow. CDH 5.1 suppresses the new exception to retain compatibility with HBase 0.96, but logs the error.
- HBASE-10339 contained code which inadvertently changed the signatures of the `getFamilyMap` method. CDH 5.1 restores these signatures to those used in HBase 0.96, to retain compatibility.

Apache Hive

- Permission inheritance fixes
- Support for decimal computation, and for reading and writing decimal-format data from and to Parquet and Avro

Hue

CDH 5.1.0 implements Hue 3.6.

New Features:

- Search App v2:
 - 100% Dynamic dashboard
 - Drag-and-Drop dashboard builder
 - Text, Timeline, Pie, Line, Bar, Map, Filters, Grid and HTML widgets
 - Solr Index creation wizard (from a file)
- Ability to view compressed Snappy, Avro and Parquet files
- Impala HA
- Close Impala and Hive sessions queries and commands

Apache Mahout

- CDH 5.1.0 implements Mahout 0.9.

See also [Apache Mahout Incompatible Changes and Limitations](#) on page 1938.

Apache Oozie

- You can now submit Sqoop jobs from the Oozie command line.
- LAST_ONLY execution mode now works correctly ([OOZIE-1319](#)).

Cloudera Search

New Features:

- A Quick Start script that automates using Search to query data from the Enron Email dataset. The script downloads the data, expands it, moves it to HDFS, indexes, and pushes the results live. The documentation now also includes a companion quick start guide, which describes the tasks the script completes, as well as customization options.
- `solrctl` now has built-in support for schema-less Solr.
- Sentry-based document-level security for role-based access control of a collection. Document-level access control associates authorization tokens with each document in the collection, enabling granting Sentry roles access to sets of documents in a collection.
- Cloudera Search includes a version of Kite 0.10.0, which includes all morphlines-related backports of all fixes and features in Kite 0.15.0. For additional information on Kite, see:
 - [Kite repository](#)
 - [Kite Release Notes](#)
 - [Kite documentation](#)
 - [Kite examples](#)
- Support for the Parquet file format is included with this version of Kite 0.10.0.
- Inclusion of `hbase-indexer-1.5.1`, a new version of the Lily HBase Indexer. This new version of the indexer includes the 0.10.0 version of Kite mentioned above. This 0.10.0 version of Kite includes the backports and fixes included in Kite 0.15.0.

Apache Sentry (incubating)

- CDH 5.1.0 implements Sentry 1.2. This includes a database-backed Sentry service which uses the more traditional GRANT/REVOKE statements instead of the previous policy file approach making it easier to maintain and modify privileges.
- Revised authorization privilege model for Hive and Impala.

Apache Spark

- CDH 5.1.0 implements Spark 1.0.
- The `spark-submit` command abstracts across the variety of deployment modes that Spark supports and takes care of assembling the classpath for you.
- Application History Server (SparkHistoryServer) improves monitoring capabilities.
- You can launch PySpark applications against YARN clusters. PySpark currently only works in YARN Client mode.

Other improvements include:

CDH 5 Release Notes

- Streaming integration with Kerberos
- Addition of more algorithms to MLlib (Sparse Vector Support)
- Improvements to Avro integration
- Spark SQL alpha release (new SQL engine). Spark SQL allows you to run SQL statements inside a Spark application that manipulate and produce RDDs.



Note:

Because of its immaturity and alpha status, Cloudera does not currently offer commercial support for Spark SQL, but bundles it with our distribution so that you can try it out.

- Authentication of all Spark communications

What's New in CDH 5.1.2

This is a maintenance release which fixes several issues. See [Issues Fixed in CDH 5.1.2](#) on page 2374



Note:

There is no CDH 5.1.1 release. This skip in the CDH 5.x sequence allows the CDH and CM components of Cloudera Enterprise 5.1.2 to have consistent numbering.

What's New in CDH 5.1.3

This is a maintenance release that fixes several issues. See [Issues Fixed in CDH 5.1.3](#) on page 2373.

What's New in CDH 5.1.4

This is a maintenance release that fixes important security issues. See [Issues Fixed in CDH 5.1.4](#) on page 2372,

What's New in CDH 5.1.5

This is a maintenance release that fixes several issues. See [Issues Fixed in CDH 5.1.5](#) on page 2371.

What's New in CDH 5.0.x

Use these links to go to a specific release.

What's New in CDH 5.0.0

This is a major release which includes new features, changes, and fixed issues. See also [Issues Fixed in CDH 5.0.x](#) on page 2376 and [Known Issues in CDH 5](#) on page 1944.

For information about CDH 5 Beta releases, see [What's New In CDH 5 Beta Releases](#) on page 1861.

New Features and Changes in CDH 5.0.0

Apache Hadoop

Apache Hadoop MapReduce Version 1 (MRv1) and Version 2 (MRv2)



Important: Cloudera recommends that you use YARN (now production-ready) with CDH 5.

- **MapReduce 2.0 (MRv2):** CDH 5 includes MapReduce 2.0 (MRv2) running on YARN. The YARN architecture splits up the two primary responsibilities of the JobTracker — resource management and job scheduling/monitoring — into separate daemons: a global ResourceManager (RM) and per-application ApplicationMasters (AM). With MRv2, the ResourceManager (RM) and per-node NodeManagers (NM) form the data-computation framework. The ResourceManager service effectively replaces the functions of the JobTracker, and NodeManagers run on worker nodes instead of TaskTracker daemons. The per-application ApplicationMaster is, in effect, a framework-specific

library and is tasked with negotiating resources from the ResourceManager and working with the NodeManagers to execute and monitor tasks. See [Apache Hadoop NextGen MapReduce \(YARN\)](#) for more information.

- **MapReduce Version 1 (MRv1):** For backward compatibility, CDH 5 continues to support the original MapReduce JobTracker and TaskTrackers, but you should migrate to MRv2.



Note:

Cloudera does not support running MRv1 and YARN daemons on the same nodes at the same time.

- **Deprecated properties:**

In Hadoop 2.0.0 and later (MRv2), several Hadoop and HDFS properties have been deprecated. (The change dates from Hadoop 0.23.1, on which the Beta releases of CDH 4 were based). See [Hadoop Deprecated Properties](#) for a list of deprecated properties and their replacements.



Note: Deprecated properties continue to work in MRv1. However, new `mapreduce*` properties do not work in MRv1.

HDFS

New Features:

- [HDFS-5776](#)- Hedged reads in HDFS for improved HBase MTTR.
- [HDFS-4685](#)- Implementation of extended file access control lists in HDFS.

Notable Bug Fixes:

- [HDFS-5339](#) - WebHDFS URI does not accept logical nameservices when security is enabled.
- [HDFS-5898](#) - Allow NFS gateway to login/relogin from its Kerberos keytab.
- [HDFS-5921](#) - "Browse filesystem" on the Namenode UI does not work if any directory has the sticky bit set.
- HDFS and Hive replication between different Kerberos realms now works.
- [HDFS-5922](#) - DataNode heartbeat thread can get stuck in a tight loop.

MapReduce & YARN

New Feature:

- FairScheduler supports moving running applications between queries.

Notable Bug Fixes:

- Several critical fixes to stabilize ResourceManager HA - Web UI, unmanaged ApplicationMasters and secure-cluster support.
- Support for large values of `mapreduce.task.io.sort.mb`.
- JobHistory Server has information on failed MapReduce jobs.

Apache HBase

New Features:

- [HBASE-10436](#)- Restore RegionServer lists removed from HBase 0.96.0 JMX.

Many of the metrics exposed in CDH 4/0.94 were removed with the refactorization of metrics in CDH 5/0.96. This patch restores the availability of the lists of live and dead RegionServers. In 0.94 this was a large nested structure as shown below, which included the RegionServer lists and metrics from each region.

```
{
  "name" : "hadoop:service=Master,name=Master",
  "modelerType" : "org.apache.hadoop.hbase.master.MXBeanImpl",
  "ZookeeperQuorum" : "localhost:2181",
```

```

....
  "RegionsInTransition" : [ ],
  "RegionServers" : [ {
    "key" : "localhost,48346,1390857257246",
    "value" : {
      "load" : 2,
    }
  }
....

```

CDH 5 Beta 1 and Beta 2 did not contain this list; they only displayed counts of the number of live and dead RegionServers. As of CDH 5.0.0, this list is now presented in a semi-colon separated field as follows:

```

{
  "name" : "Hadoop:service=HBase,name=Master,sub=Server",
  "modelerType" : "Master,sub=Server",
  "tag.Context" : "master",
  "tag.liveRegionServers" : "localhost,56196,1391992019130",
  "tag.deadRegionServers" :
"localhost,40010,1391035309673;localhost,41408,1391990380724;localhost,38682,1390950017735",
  ...
}

```

- Assorted usability and compatibility improvements as well as improvements to exporting snapshots.

Apache Flume

New Feature:

- The HBase Sink now supports coalescing multiple Increment RPCs into one ([FLUME-2338](#)).

Changed Behavior:

- File Channel Write timeout has been removed and the configuration parameter is now ignored ([FLUME-2307](#)).
- Syslog UDP source can now accept larger messages ([FLUME-2130](#)).
- AsyncHBase Sink is now fully functional ([FLUME-2334](#)).
- Use standard lookup to find queue/topic in JMS Source ([FLUME-2311](#)).

Notable Bug Fixes:

- Deadlock fixed in Dataset sink ([FLUME-2320](#)).
- FileChannel Dual Checkpoint Backup Thread is now released on application stop ([FLUME-2328](#)).
- Spool Dir source now checks interrupt flag before writing to channel ([FLUME-2283](#)).
- Morphline sink increments `eventDrainAttemptCount` when it takes event from channel ([FLUME-2323](#)).
- Bucketwriter now permanently closed only on idle and roll timeouts ([FLUME-2325](#)).
- `BucketWriter#close` now cancels `idleFuture` ([FLUME-2305](#)).

Apache Oozie

As of CDH 5.0.0 Oozie includes a glob pattern feature ([OOZIE-1471](#)), allowing you do a move of wild cards in the FS Action. For example:

```

<fs name="archive-files">
<move source="hdfs://namenode/output/*"
target="hdfs://namenode/archive" />
<ok to="next"/>
<error to="fail"/>
</fs>

```

By default, up to 1000 files can be matched; you can change this default by means of the `oozie.action.fs.glob.max` parameter.

Cloudera Search

- Cloudera Search includes a version of Kite 0.10.0, which includes backports of all fixes and features in Kite 0.12.0. For additional information on Kite, see:

- [Kite repository](#)
- [Kite Release Notes](#)
- [Kite documentation](#)
- [Kite examples](#)

What's New in CDH 5.0.1

This is a maintenance release which fixes several issues. In addition, it introduces a change to the configuration for HTTPS communication between HDFS and YARN. See also [Issues Fixed in CDH 5.0.1](#) on page 2379.

New Features and Changes in CDH 5.0.1

Enabling TLS/SSL in CDH 5: Enabling HTTPS communication in CDH 5 requires extra configuration properties to be added to YARN (`yarn-site.xml` and `mapred-site.xml`) and HDFS (`hdfs-site.xml`), in addition to the existing configuration settings described [here](#).

What's New in CDH 5.0.2

This is a maintenance release which fixes several issues. See [Issues Fixed in CDH 5.0.2](#) on page 2378.

What's New in CDH 5.0.3

This is a maintenance release that fixes several issues. See [Issues Fixed in CDH 5.0.3](#) on page 2377.

What's New in CDH 5.0.4

This is a maintenance release that fixes several issues. See [Issues Fixed in CDH 5.0.4](#) on page 2377.

What's New in CDH 5.0.5

This is a maintenance release that fixes the “POODLE” and Apache Hadoop Distributed Cache vulnerabilities described in [“POODLE” Vulnerability on TLS/SSL enabled ports](#) on page 2376, as well as other issues. All CDH 5.0.x users should upgrade to 5.0.5 as soon as possible. See [Issues Fixed in CDH 5.0.5](#) on page 2376.

What's New in CDH 5.0.6

This is a maintenance release that fixes several issues. See [Issues Fixed in CDH 5.0.6](#) on page 2376.

What's New In CDH 5 Beta Releases

Use these links to go to a specific release.

What's New in CDH 5 Beta 1

This is a beta release which previews new features, changes, and fixed issues.

Oracle JDK 7 Support

- CDH 5 supports Oracle JDK 1.7 and supports users running applications compiled with JDK 1.7. For CDH 5 Beta 1 the certified version is JDK 1.7.0_25. Cloudera has tested this version across all components.
- CDH 5 does not support JDK 1.6; you must install JDK 1.7.

Apache Flume

New Features:

- [FLUME-2190](#) - Includes a new Twitter Source that feeds off the Twitter firehose
- [FLUME-2109](#) - HTTP Source now supports HTTPS
- Flume now auto-detects Cloudera Search dependencies.

Apache Hadoop

HDFS

New Features:

- [HDFS-4953](#): Enable HDFS local reads via mmap.
- [HDFS-2802](#): Support for RW/RO snapshots in HDFS. See:
`hadoop-hdfs-project/hadoop-hdfs/src/site/apt/HdfsNfsGateway.appt.vm`
- [HDFS-4750](#): Support NFSv3 interface to HDFS.
- [HDFS-4817](#): Make HDFS advisory caching configurable on a per-file basis.
- [HDFS-3601](#): Add BlockPlacementPolicyWithNodeGroup to support block placement with 4-layer network topology.
- [HDFS-5122](#): Support failover and retry in WebHdfsFileSystem for NN HA.
- [HDFS-4772](#) / [HDFS-5043](#): Add number of children (of a directory) in HdfsFileStatus.
- [HDFS-4434](#): Provide a mapping from INodeId to INode. See: `/.reserved/.inodes/<INODE_NUMBER>`
- [HDFS-2576](#): Enhances the DistributedFileSystem's Create API so that clients can specify favored DataNodes for a file's blocks.

Changed Features:

- [HDFS-4659](#): Support setting execution bit for regular files.
 - **Impact:** In CDH 5, files copied out of `copyToLocal` may now have the executable bit set if it was set when they were created or copied into HDFS.
- [HDFS-4594](#): WebHDFS open sets Content-Length header to what is specified by length parameter rather than how much data is actually returned.
 - **Impact:** In CDH 5, Content-Length header will contain the number of bytes actually returned, rather than the request length.

Changed Behavior:

- [HDFS-4645](#): Move from randomly generated block ID to sequentially generated block ID.
- [HDFS-4451](#): HDFS balancer command returns exit code 1 on success instead of 0.

MapReduce v2 (YARN)

New Features:

- **ResourceManager High Availability:** YARN now allows you to use multiple ResourceManagers so that there is no single point of failure. In-flight jobs are recovered without re-running completed tasks.
- Monitoring and enforcing memory and CPU-based resource utilization using `cgroups`.
- **Continuous Scheduling:** This feature decouples scheduling from the node heartbeats for improved performance in large clusters.

Changed Feature:

- **ResourceManager Restart:** Persistent implementations of the RMStateStore (filesystem-based and ZooKeeper-based) allow recovery of in-flight jobs.

Apache HBase

Summary of New Features

- Support for Hadoop 2.0
- Improved MTTR (meta first recovery, distributed log replay)
- Improved compatibility and upgradeability (ProtoBuf serialization format)
- Namespaces added for administrative domains
- Snapshots (ported to 0.94 / CDH4.2)
- Online region merge mechanisms added
- Major security and functional improvements made for the REST proxy server

Administrative Features

ProtoBuf: All of the serialization that goes across the wire between servers written to and read by HBase file formats have been converted to extensible Protobuf encodings. This breaks compatibility with previous versions but should make future extensions less likely to break compatibility in these areas. This feature is enabled by default.

- [HBASE-5305](#): Improve cross-version compatibility and upgradeability.
- [HBASE-7898](#): Serializing cells over RPC.

Namespaces: Namespaces is a new feature that groups tables into different administrative domains. An admin can be only given rights to act upon a particular namespace. This feature is enabled by default and requires file system layout changes that must be completed during upgrade.

- [HBASE-8015](#): Added support for namespaces.

MTTR Improvements: Mean time to recovery has greatly improved.

- [HBASE-7590](#): “Costless” notifications from master to rs/clients.
- [HBASE-7213](#) / [HBASE-8631](#): New `.meta` suffix to separate HLog file / Recover Meta before other regions in case of server crash.
- [HBASE-7006](#): Distributed log replay (Caveat).
- [HBASE-9116](#): Adds a view/edit tool for favored node mappings for regions (incomplete, likely a dot version).

Metrics: There are several new metrics and a new naming convention for metrics in HBase. This also includes metrics for each region.

- [HBASE-3614](#): Per region metrics.
- [HBASE-4050](#): Rationalize metrics; Update HBase metrics framework to metrics2.

Miscellaneous:

- [HBASE-7403](#): HBase online region merge.
- Shell improvements; tables list to be more well-rounded.
- [HBASE-5953](#): Expose the current state of the balancerSwitch.
- [HBASE-5934](#): Add the ability for Performance Evaluation to set table compression.
- [HBASE-6135](#): New Web UI.
- [HBASE-8148](#): Allow IPC to bind to a specific address (also 0.94.7)
- [HBASE-5498](#): Secure Bulk Load (also 0.94.5)

Backup and Disaster Recovery Features

Replication: Several critical bug fixes.

- [HBASE-9373](#): Replication has been hardened.
- [HBASE-9158](#): Serious bug in cyclic replication.
- [HBASE-8737](#): Changes to the replication RPC to use cell blocks.

Snapshots: HBase table snapshots were backported to 0.94.x. There are some incompatibilities between the implementation released in CDH 4 with that in CDH 5.

- [HBASE-7290](#): Online snapshots (backported to 0.94.x).
- [HBASE-8352](#): Rename snapshots folder from `.snapshot` to `.hbase-snapshots` (Incompatible change).

Copy table:

- [HBASE-8609](#): Add startRow-stopRow options to the CopyTable.

Import:

- [HBASE-7702](#): Add filtering to import jobs.

HBase Proxies

The REST server now supports Hadoop authentication and authorization mechanisms. The Avro gateway has been removed while the Thrift2 proxy has made progress but is not complete. However, it has been included as a preview feature.

REST:

- [HBASE-9347](#): Support for specifying filter in REST server requests.
- [HBASE-7803](#): Support caching on scan.
- [HBASE-7757](#): Add Web UI for Thrift and REST servers.
- [HBASE-5050](#): SPNEGO-based authentication.
- [HBASE-8661](#): Support REST over HTTPS.
- [HBASE-8662](#): Support for impersonation.
- [HBASE-7986](#): [REST] Make HTablePool size configurable.

Thrift:

- [HBASE-5879](#): Enable JMX metrics collection for the Thrift proxy.

Thrift2: Ongoing efforts to match Thrift and REST functionality. (Incomplete, only a preview feature)

Avro:

- [HBASE-5948](#): Avro gateway removed.

Stability Features

There have been several bug fixes, test fixes and configuration default changes that greatly increase our confidence in the stability of the 0.96.0 release. The main improvement comes from the use of a systematic fault-injection framework.

- [HBASE-7721](#) Atomic multi-row mutations in META
- Integration testing
- [HBASE-7977](#) - TableLocks
- [HBASE-7898](#) [many flaky tests hardened](#)

Performance Features

Several features have been added to improve throughput and performance characteristics of HBase and its clients.



Warning:

Currently the 0.95.2/CDH 5 beta 1 release will suffer performance degradation when over 40 nodes are used when compared to CDH 4.

Throughput:

- [HBASE-4676](#): Prefix compression / tree encoding.
- [HBASE-8334](#): Essential column families on by default (filtering optimization).
- [HBASE-5074](#) / [HBASE-8322](#): Re-enable HBase checksums by default.
- [HBASE-6466](#): Enable multi-threaded memstore flush
- [HBASE-6783](#): Make short circuit read the default.

Predictable Performance:

- [HBASE-5959](#): Added a Stochastic LoadBalancer
- [HBASE-7842](#): Exploring compactor.
- [HBASE-7236](#): Add per-table/per-cf configuration via metadata
- [HBASE-8163](#): MemStoreChunkPool: Improvement for Java GC

- [HBASE-4391](#)[HBASE-6567](#): Mlock / Memory locking improvements (less disk swap).
- [HBASE-4391](#): Bucket cache (untested)

Miscellaneous:

- [HBASE-6870](#): Improvement to HTable coprocessorExec scan performance.

Developer Features

These features are to aid application developers or for major changes that will enable future minor version improvements.

- [HBASE-9121](#): HTrace updates.
- [HBASE-8375](#): Durability setting per table.
 - [HBASE-7801](#) Deferred sync for WAL logs (0.94.7 and later)
- [HBASE-7897](#): Tags supported in cell interface (for future security features).
- [HBASE-5937](#): Refactor HLog into interface (allows for new HLogs in 0.96.x).
- [HBASE-4336](#): Modularization of POM / Multiple jars (many follow-ons, [HBASE-7898](#)).
- [HBASE-8224](#): Publish -hadoop1 and -hadoop2 versioned jars to Maven (CDH published jars are assumed -hadoop2).
- [HBASE-9164](#): Move towards Cell interface in client instead of KeyValue.
- [HBASE-7898](#): Serializing cells over RPC.
- [HBASE-7725](#): Add ability to create custom compaction request.

Hue

New Features:

- With the Sqoop 2 application, data from databases can be easily exported or imported into HDFS in a scalable manner. The Job Wizard hides the complexity of creating Sqoop jobs and the dashboard offers live progress and log access.
- Zookeeper App: Navigate and browse the Znode hierarchy and content of a Zookeeper cluster. Znodes can be added, deleted and edited. Multi-clusters are supported and various statistics are available for them.
- The Hue Shell application has been removed and replaced by the Pig Editor, HBase Browser and the Sqoop 1 apps.
- Python 2.6 is required.
- Beeswax daemon has been replaced by HiveServer2.
- CDH 5 Hue will only work with HiveServer2 from CDH 5. No support for impersonation.

Hue also includes the following changed features (Updated to upstream version 3.0.0):

- [\[HUE-897\]](#) - [core] Redesign of the overall layout
- [\[HUE-1521\]](#) - [core] Improve JobTracker High Availability
- [\[HUE-1493\]](#) - [beeswax] Replace the Beeswax server with HiveServer2
- [\[HUE-1474\]](#) - [core] Upgrade Django backend version from 1.2 to 1.4
- [\[HUE-1506\]](#)- [search] Impersonation support added
- [\[HUE-1475\]](#) - [core] Switch back from the Spawning web server
- [\[HUE-917\]](#) - Support SAML based authentication to enable single sign-on (SSO)

From master:

- [\[HUE-950\]](#) - [core] Improvements to the document model
- [\[HUE-1595\]](#) - Integrate Metastore data into Hive and Impala Query UIs
- [\[HUE-1275\]](#) - [metastore] Show Metastore table details
- [\[HUE-1622\]](#) - [core] Mini tour added to Hue home page

Apache Hive and HCatalog

New Features (Updated to upstream version 0.11.0):

- [\[HIVE-446\]](#) - Implement TRUNCATE for table data
- [\[HIVE-896\]](#) - Add LEAD/LAG/FIRST/LAST analytical windowing functions to Hive

- [\[HIVE-2693\]](#) - Add DECIMAL data type
- [\[HIVE-3834\]](#) - Support ALTER VIEW AS SELECT in Hive

Performance improvements (from 0.12):

- [\[HIVE-3764\]](#) - Support metastore version consistency check
- [\[HIVE-305\]](#) - Port Hadoop streaming process's counters/status reporters to Hive Transforms
- [\[HIVE-1402\]](#) - Add parallel ORDER BY to Hive
- [\[HIVE-2206\]](#) - Add a new optimizer for query correlation discovery and optimization
- [\[HIVE-2517\]](#) - Support GROUP BY on struct type
- [\[HIVE-2655\]](#) - Ability to define functions in HQL
- [\[HIVE-4911\]](#) - Enable QOP configuration for HiveServer2 Thrift transport

Cloudera Impala

Cloudera Impala 1.2.0 is now available as part of CDH 5. For more details on Impala, refer the [Impala Documentation](#).

Llama

Llama is a system that mediates resource management between Cloudera Impala and Hadoop YARN. Llama enables Impala to reserve, use, and release resource allocations in a Hadoop cluster. Llama is only required if resource management is enabled in Impala.

Apache Mahout

New Features (Updated to Mahout 0.8):

- Numerous performance improvements to Vector and Matrix implementations, APIs and their iterators (see also [MAHOUT-1192](#), [MAHOUT-1202](#))
- Numerous performance improvements to the recommender implementations (see also [MAHOUT-1272](#), [MAHOUT-1035](#), [MAHOUT-1042](#), [MAHOUT-1151](#), [MAHOUT-1166](#), [MAHOUT-1167](#), [MAHOUT-1169](#), [MAHOUT-1205](#), [MAHOUT-1264](#))
- [MAHOUT-1088](#): Support for biased item-based recommender.
- [MAHOUT-1089](#): SGD matrix factorization for rating prediction with user and item biases.
- [MAHOUT-1106](#): Support for SVD++
- [MAHOUT-944](#): Support for converting one or more Lucene storage indexes to SequenceFiles as well as an upgrade of the supported Lucene version to Lucene 4.3.
- [MAHOUT-1154](#) and related: New streaming k-means implementation that offers online (and fast) clustering.
- [MAHOUT-833](#): Make conversion to SequenceFiles Map-Reduce. 'seqdirectory' can now be run as a MapReduce job.
- [MAHOUT-1052](#): Add an option to MinHashDriver that specifies the dimension of vector to hash (indexes or values).
- [MAHOUT-884](#): Matrix concatenate utility; presently only concatenates two matrices.

Apache Oozie

New Features:

- Updated to Oozie 4.0.0.
- **High Availability:** Multiple Oozie servers can now be utilized to provide an HA Oozie service as well as provide horizontal scalability. See upstream [documentation](#) for more details.
- **HCatalog Integration:** HCatalog table partitions can now be used as data dependencies in coordinators. See upstream [documentation](#) for more details. .
- **SLA Monitoring:** Oozie can now actively monitor SLA-sensitive jobs and send out notifications for SLA meets and misses. SLA information is also now available through a new SLA tab in the Oozie Web UI, JMS messages, and a REST API. See upstream [documentation](#).
- **JMS Notifications:** Oozie can now publish notifications to a JMS Provider about job status changes and SLA events. See upstream [documentation](#).
- The FileSystem action can now use glob patterns for file paths when doing `move`, `delete`, `chmod`, and `chgrp`.

Cloudera Search

Cloudera Search 1.0.0 is now available as part of CDH 5. For more details on Search see the [Search documentation](#).

The Cloudera Development Kit (CDK) is a set of libraries and tools that can be used with Search and other CDH components to build jobs/systems on top of the Hadoop ecosystem. See the [CDK Documentation](#) and [Release Notes](#) for more details.



Note: An existing dependency, Apache Tika, has been upgraded to version 1.4.

Apache Sentry (incubating)

CDH 5 Beta 1 includes the first upstream release of Apache Sentry, `sentry-1.2.0-incubating`.

Apache Sqoop

CDH 5 Sqoop 1 has been rebased on Apache Sqoop 1.4.4.

What's New in CDH 5 Beta 2

This is a beta release which previews new features, changes, and fixed issues. See also [Issues Fixed in CDH 5 Beta 2](#) on page 2383.

New Features and Changes in CDH 5 Beta 2

Apache Crunch

The Apache Crunch™ project develops and supports Java APIs that simplify the process of creating data pipelines on top of Apache Hadoop. The Crunch APIs are modeled after FlumeJava (PDF), which is the library that Google uses for building data pipelines on top of their own implementation of MapReduce.

Apache DataFu

- Upgraded from version 0.4 to 1.1.0 (this upgrade is not backward compatible).
- New features include UDFS SHA, `SimpleRandomSample`, `COALESCE`, `ReservoirSample`, `EmptyBagToNullFields`, and many others.

Apache Flume

- [FLUME-2294](#) - Added a new sink to write Kite datasets.
- [FLUME-2056](#) - Spooling Directory Source can now only pass the name of the file in the event headers.
- [FLUME-2155](#) - File Channel is indexed during replay to improve replay performance for faster startup.
- [FLUME-2217](#) - Syslog Sources can optionally preserve all syslog headers in the message body.
- [FLUME-2052](#) - Spooling Directory Source can now replace or ignore malformed characters in input files.

Apache Hadoop

HDFS

New Features/Improvements:

- As of CDH 5 Beta 2, you can upgrade HDFS with high availability (HA) enabled, if you are using Quorum-based storage. (Quorum-based storage is the only method available in CDH 5; NFS shared storage is not supported.)
- [HDFS-4949](#) - CDH 5 Beta 2 supports HDFS caching.
- As of CDH 5 Beta 2, you can configure an NFSv3 gateway that allows any NFSv3-compatible client to mount HDFS as a file system on the client's local file system.
- [HDFS-5709](#) - Improve upgrade with existing files and directories named `.snapshot`.

Major Bug Fixes:

- [HDFS-5449](#) - Fix WebHDFS compatibility break.
- [HDFS-5671](#) - Fix socket leak in `DFSInputStream#getBlockReader`.
- [HDFS-5353](#) - Short circuit reads fail when `dfs.encrypt.data.transfer` is enabled.
- [HDFS-5438](#) - Flaws in block report processing can cause data loss.

Changed Behavior:

- As of CDH 5 Beta 2, in order for the NameNode to start up on a secure cluster, you should have the `dfs.web.authentication.kerberos.principal` property defined in `hdfs-site.xml`. For clusters managed by Cloudera Manager, you do not need to explicitly define this property.
- [HDFS-5037](#) - Active NameNode should trigger its own edit log rolls. Clients will now retry for a configurable period when encountering a NameNode in Safe Mode.
- The default behavior of the `mkdir` command has changed. As of CDH 5 Beta 2, if the parent folder does not exist, the `-p` switch must be explicitly mentioned otherwise the command fails.

MapReduce (MRv1 and YARN)

- Fair Scheduler (in YARN and MRv1) now supports advance configuration to automatically place applications in queues.
- MapReduce now supports running multiple reducers in `uber` mode and in local job runner.

Apache HBase

- **Online Schema Change** is now a supported feature.
- **Online Region Merge** is now a supported feature.
- **Namespaces:** CDH 5 Beta 2 includes the namespaces feature which enables different sets of tables to be administered by different administrative users. All upgraded tables will live in the default "hbase" namespace. Administrators may create new namespaces and create tables users with rights to the namespace may administer permissions on the tables within the namespace.
- There have been several improvements to HBase's **mean time to recovery** (mttr) in the face of Master or RegionServer failures.
 - Distributed log splitting has matured, and is always activated. The option to use the old slower splitting mechanism no longer exists.
 - Failure detection time has been improved. New notifications are now sent when RegionServers or Masters fail which triggers corrective action quickly.
 - The Meta table has a dedicated write ahead log which enables faster recovery region recovery if the RegionServer serving meta goes down.
- The **Region Balancer** has been significantly updated to take more load attributes into account.
- Added **TableSnapshotInputFormat** and **TableSnapshotScanner** to perform scans over HBase table snapshots from the client side, bypassing the HBase servers. The former configures a MapReduce job, while the latter does a single client-side scan over snapshot files. Can also be used with offline HBase with in-place or exported snapshot files.
- The KeyValue API has been deprecated for applications in favor of the **Cell** interface. Users upgrading to HBase 0.96 may still use KeyValue by future upgrades may remove the class or parts of its functionality. Users are encouraged to update their applications to use the new Cell interface.
- Currently Experimental features:
 - **Distributed log replay:** This mechanism allows for faster recovery from RegionServer failures but has one special case where it will violate ACID guarantees. Cloudera does not currently recommend activating this feature.
 - **Bucket cache:** This is an offheap caching mechanism that use extra RAM and block devices (such as flash drives) to greatly increase the read caching capabilities provided by the BlockCache. Cloudera does not currently recommend activating this feature.
 - **Favored nodes:** This feature enables HBase to better control where its data is written to in HDFS in order to better preserve performance after a failure. This is disabled currently because it doesn't interact well with the HBase Balancer or HDFS Balancer. Cloudera does not currently recommend activating this feature.

See this [blog post](#) for more details.

*Apache Hive***New Features:**

- Improved JDBC specification coverage:

- Improvements to `getDatabaseMajorVersion()`, `getDatabaseMinorVersion()` APIs ([HIVE-3181](#))
- Added JDBC support for new datatypes: Char ([HIVE-5683](#)), Decimal ([HIVE-5355](#)) and Varchar ([HIVE-5209](#))
- You can now specify the database for a session in the HiveServer2 connection URL ([HIVE-4256](#))
- Encrypted communication between the Hive Server and Clients. This includes TLS/SSL encryption for non-Kerberos connections to HiveServer2 ([HIVE-5351](#)).
- A native Parquet SerDe is now available as part of the CDH 5 Beta 2 package. Users can directly create a Parquet format table without any external package dependency.

Changed Behavior:

- [HIVE-4256](#) - With Sentry enabled, the `use <database>` command is now executed as part of the connection to HiveServer2. Hence, a user with no privileges to access a database will not be allowed to connect to HiveServer2.

Hue

- Hue has been upgraded to version 3.5.0.
- Impala and Hive Editor are now one-page apps. The Editor, Progress, Table list and Results are all on the same page
- Result graphing for the Hive and Impala Editors.
- Editor and Dashboard for Oozie SLA, crontab and credentials.
- The Sqoop2 app supports autocomplete of database and table names/fields.
- [DBQuery App](#): MySQL and PostgreSQL Query Editors.
- New Search feature: [Graphical facets](#)
- Integrate external Web applications in any language. See this [blog post](#) for more details.
- Create Hive tables and load quoted CSV data. Tutorial available [here](#).
- Submit any Oozie jobs directly from HDFS. Tutorial available [here](#)
- New [SAML backend](#) enables single sign-on (SSO) with Hue.

Apache Oozie

- Oozie now supports cron-style scheduling capability.
- Oozie now supports High Availability with security.

Apache Pig

- `AvroStorage` rewritten for better performance, and moved from `piggybank` to core Pig
- `ASSERT`, `IN`, and `CASE` operators added
- `ParquetStorage` added for integration with Parquet

Cloudera Search

- The Cloudera CDK has been renamed and updated to Kite version 0.11.0. For additional information on Kite, see:
 - [Kite repository](#)
 - [Kite Release Notes](#)
 - [Kite documentation](#)
 - [Kite examples](#)

Apache Spark (incubating)

Spark is a fast, general engine for large-scale data processing. For installation and configuration instructions, see [Setting Up Apache Spark Using the Command Line](#).

Apache Sqoop

Sqoop 2 has been upgraded from version 1.99.2 to 1.99.3.

What's New in Apache Impala

This release of Impala contains the following changes and enhancements from previous releases.

New Features in CDH 5.16.x

The following are some of the most significant new features in this release.

Fine Grained Privileges

Sentry and Impala introduced fine grained privileges to provide object-level privileges to roles.

Fine grained privileges include the `REFRESH` and `CREATE` privileges, which allow users to create databases and tables, and to execute commands that update metadata information on Impala databases and tables. See [Impala Sentry documentation](#) for the new privileges and the scopes of the objects that you can grant the new privileges on.

The following new privileges were added:

- The `REFRESH` privilege
- The `CREATE` privilege
- The `SELECT` and `INSERT` privileges on `SERVER`

If a role has `SELECT` or `INSERT` privilege on an object in Impala before upgrading to CDH 5.16.1, that role will automatically get the `REFRESH` privilege during the upgrade.

Object Ownership

Object ownership designates an owner for a database, table, or view in Sentry. The owner of an object has the `OWNER` privilege which is the equivalent of the `ALL` privilege on the object. See the [Object Ownership](#) documentation for information about enabling object ownership.

If the object ownership feature is enabled, Sentry grants the user the `OWNER` privilege. Whether or not object ownership is enabled, HMS stores the object creator as the default object owner. Previously, HMS stored the Kerberos user as the object owner.

The following statements were added to Impala to support object ownership via Sentry:

- `ALTER DATABASE SET OWNER`
- `ALTER TABLE SET OWNER`
- `ALTER VIEW SET OWNER`
- `SHOW GRANT USER`

Admission Control Enhancement

A new query option, [MAX_MEM_ESTIMATE_FOR_ADMISSION](#), was added. Use the new option to set an upper limit on the memory estimates of a query as a workaround for over-estimates precluding a query from being admitted.

General Performance Improvements

A new query option, [SHUFFLE_DISTINCT_EXPRS](#), controls the shuffling behavior when a query has both grouping and distinct expressions.

Metadata Performance Improvements

- Incremental Stats

The following enhancements improve Impala stability. The features reduce chances of having `catalogd` and `impalad` crash due to be out of memory when using incremental stats.

- Incremental stats are now compressed in memory in `catalogd`, reducing memory footprint in `catalogd`.
- Incremental stats are fetched on demand from `catalogd` by `impalad` coordinators. This enhancement reduces memory footprint for `impalad` coordinators and `statestored` and also reduces network requirements to broadcast metadata.

See [Loading Incremental Statistics from Catalogd](#) for details.

- Automatic Invalidation of Metadata



Note: This feature is experimental and not recommended for use in production clusters.

To keep the size of metadata bounded and to reduce the chances of `catalogd` cache running out of memory, this release introduces an automatic metadata invalidation feature with time-based and memory-based invalidation.

Automatic invalidation of metadata provides more stability with lower chances of running out of memory, but could potentially cause performance risks. The feature is turned off by default.

See [Startup Options for Automatic Invalidation of Metadata](#) for details.

Compatibility and Usability Enhancements

- Additional separators are supported between date and time in default `TIMESTAMP` format, specifically, the multi-space separator and the 'T' separator. See [TIMESTAMP Data Type](#) for more information on `TIMESTAMP` format.
- New hint placement is supported for `INSERT` statements. See [Optimizer Hints in Impala](#) for detail.
- The `REGEX_ESCAPE()` function was implemented for escaping special characters to treat them literally in string literals.
- `SHOW CREATE VIEW` was implemented with the same functionality as `SHOW CREATE TABLE`.
- Impala Shell can connect directly to `impalad` when configured with proxy load balancer and Kerberos. See [Configuration Options](#) for the new flag that enables the direct connection.

New Features in CDH 5.15.x

The following are some of the most significant new features in this release:

- Added the `TABLESAMPLE` clause in the `COMPUTE STATS` statement. See [TABLESAMPLE](#) for the new `TABLESAMPLE` clause.
- Extended `COMPUTE STATS` to support a list of columns. See [COMPUTE STATS](#) for the new syntax.
- Added the new `COMPUTE_STATS_MIN_SAMPLE_SIZE` query option. The query option specifies the minimum number of bytes that will be scanned in `COMPUTE STATS TABLESAMPLE`, regardless of the user-supplied sampling percent. See [COMPUTE_STATS_MIN_SAMPLE_SIZE](#).
- Added a `TBLPROPERTY` for controlling stats extrapolation on a per-table basis: `impala.enable.stats.extrapolation=true/false`. See [stats extrapolation](#) for information about stats extrapolation.
- Added the new built-in `regex_escape` function. The function instructs Impala to interpret the following special characters literally rather than as special characters: `.\+*?[^\]${}=!<>|:-`
See [regex_escape](#) for information about the `regex_escape` function.
- Enhanced the existing `ltrim` and `rtrim` string functions to accept an argument that specifies a set of characters to be trimmed from the input string. See [ltrim and rtrim](#) for information about the functions.
- Implemented `murmur_hash` function. See [murmur_hash](#) for information about the new function.
- Introduced the support for the Kudu `DECIMAL` type in Kudu 1.7.0.
- Now Impala maps a signed integer logical type in Parquet to a supported Impala column type as below:
 - `INT_8` -> `TINYINT`
 - `INT_16` -> `SMALLINT`
 - `INT_32` -> `INT`
 - `INT_64` -> `BIGINT`
- Parquet dictionary filtering now works on nested data.
- Based on the existing Parquet column chunk level statistics `null_count`, Impala's Parquet scanner was enhanced to skip an entire row group if the `null_count` statistics indicate that all the values under the predicated column are `NULL` as no result rows would be returned from that row group.

- The Oracle-style hint placement for `INSERT` statements is now supported. See for information on hints in Impala SQL.
- Insert plan hints for `CREATE TABLE AS SELECT` are now supported. See for information on hints in Impala SQL.
- Improved concurrency of DDL and DML operations during catalog updates.
- The statestore update logic was improved to reduce issues, such as too many queries being admitted by different coordinators, or queries being queued for longer than necessary and blocking subsequent updates to different topics.
- The size limit for statestore updates was increased, and copying of the metadata and reduce the memory footprint were reduced. Now the catalog objects are passed and (de)compressed between FE and BE one at a time.

New Features in CDH 5.14.x

The following are some of the most significant new features in this release:

- **New Boolean operators:** `IS TRUE`, `IS NOT TRUE`, `IS FALSE`, `IS NOT FALSE`, `IS UNKNOWN`, and `IS NOT UNKNOWN`. These operators always return `TRUE` or `FALSE`, even if the left-hand expression evaluates to `NULL`. `IS [NOT] TRUE` and `IS [NOT] FALSE` are `NULL`-safe equivalents for equality and inequality comparisons. `IS [NOT] UNKNOWN` is identical to `IS [NOT] NULL`.
- The new built-in function `DATE_TRUNC()` lets you truncate a `TIMESTAMP` value to a given precision. For example, calling `date_trunc('hour', ts)` truncates `ts` to the beginning of the corresponding hour, with all minutes, seconds, milliseconds, and so on set to zero.
- The `impala-shell` command now accepts a command-line argument `--query_option=option=value` to set a query option for that shell session. You can specify multiple query options by supplying multiple instances of `--query_option` in the `impala-shell` invocation.

Configuration files for `impala-shell`, such as `.impalarc`, can also contain a section beginning with `[impala.query_options]`, followed by lines of the form `query_option=value`. These query options take effect by default for all `impala-shell` sessions that use the specified configuration file.

- The `SET` statement, when run without any argument, now divides its output into groups corresponding to “regular”, “advanced”, “development” (unfinished, not ready for general use), and “deprecated” (subject to removal). By default, the “development” and “deprecated” option groups are not shown in `SET` output. Issue the statement `SET ALL` to see the options and values from all groups.
- Impala optimizes some join queries against Kudu tables by constructing runtime filters representing the minimum and maximum matching values from the join columns. The min-max information is “pushed down” to Kudu so that Kudu can prune unnecessary rows before sending the result set back to Impala.
- The allowed range for the `BATCH_SIZE` query option is now 0 to 65536. Formerly, the upper bound was undefined.
- Arithmetic expressions involving both `DECIMAL` and `FLOAT` or `DOUBLE` arguments now produce results of `DECIMAL` type instead of `DOUBLE`.

New Features in CDH 5.13.x

The following are some of the most significant new features in this release:

- Improvements to memory management through the use of a buffer pool. This mechanism allows queries to use less memory, reserves the required memory during query startup, and reduces the frequency of out-of-memory errors. It makes query planning and memory estimation more accurate, so that if a query begins executing it is unlikely to encounter an out-of-memory error partway through. The memory buffer used during spill-to-disk processing is smaller: instead of 8 MiB, this buffer defaults to 2 MiB and Impala can reduce it to as little as 64 KiB when appropriate.

This feature includes new query options for fine-tuning memory areas used during query processing: `MIN_SPILLABLE_BUFFER_SIZE`, `DEFAULT_SPILLABLE_BUFFER_SIZE`, `MAX_ROW_SIZE`, and `BUFFER_POOL_LIMIT`.

- Improvements to the mechanism for caching HDFS file handles. This caching mechanism improves the performance and scalability of queries that access the same file multiple times, for example to retrieve different columns from a Parquet file. Caching the file handle across `open()` calls reduces the load on the HDFS NameNode.


This feature is currently disabled by default. It is enabled by setting a non-zero value for the `max_cached_file_handles` configuration setting. Currently, ETL processes that append to existing HDFS files or overwrite HDFS files in place can interact with this setting in a way that turns off short-circuit reads for some `impalad` hosts. See [HDFS-12528](#) for tracking information.

- A new command in `impala-shell`, `rerun` or its abbreviation `@`, lets you re-execute previous commands based on their numbering in the `history` output.
- You can specify the minimum required TLS/SSL version using the `--ssl_minimum_version` setting, for example `--ssl_minimum_version=tlsv1.2`.
- You can specify the set of allowed TLS ciphers using the `--ssl_cipher_list` configuration setting. See the output of `man ciphers` for the full set of keywords and notation allowed in the argument string.
- New or enhanced built-in functions:
 - `trunc()` can now apply to numeric types (`FLOAT`, `DOUBLE`, and `DECIMAL`) in addition to `TIMESTAMP`. Although this functionality was already available through the `truncate()` function, the new signatures for `trunc()` make it easier to port code from other popular database systems to Impala.
 - A new date/time function `utc_timestamp()` provides a simple way to get a stable, interoperable representation of a `TIMESTAMP` value without using a chain of functions to convert between representations and apply a specific timezone.
- The `CREATE TABLE LIKE PARQUET` statement can now handle Parquet files produced outside of Impala and containing `ENUM` types. The `ENUM` columns become `STRING` columns in the target table, and the `ENUM` values are turned into corresponding `STRING` values.
- Kudu enhancements:
 - You can now create a Kudu table without using a `PARTITION BY` clause. Kudu automatically creates a single partition to cover the entire possible range of values. This feature is intended for small lookup tables, where the overhead of partitioning does not make sense for tables that typically have a full table scan for each query.
 - More granular Sentry authorization for Kudu tables. Kudu tables can now use column-level privileges. The `SELECT` and `INSERT` statements can now use the corresponding `SELECT` and `INSERT` privileges. Other Kudu statements still require the `ALL` privilege.
 - The `ALTER TABLE` statement can modify a number of storage attributes for the columns of Kudu tables. You can use the `ALTER COLUMN` clause of `ALTER TABLE` along with the `SET` keyword to change the properties `DEFAULT`, `BLOCK_SIZE`, `ENCODING`, and `COMPRESSION`. You can use the `DROP DEFAULT` clause to remove the default value from a column.
- For non-Kudu tables, you can use the `ALTER TABLE` syntax `ALTER COLUMN col SET COMMENT 'text'` to change the comment for an individual column.

New Features in CDH 5.12.x

The following are some of the most significant new features in this release:

- Impala can now read and write data stored on the Microsoft Azure Data Lake Store (ADLS). See [Using Impala with the Azure Data Lake Store \(ADLS\)](#) for details.
- New built-in functions:
 - A new string function, `replace()`, which is faster than `regexp_replace()` for simple string substitutions. See [Impala String Functions](#) for details.

- A new conditional function, `nv12()`, which offers more flexibility than the `nv1()` function. It lets you return one value for `NOT NULL` arguments, and a different value for `NULL` arguments. See [Impala Conditional Functions](#) for details.
 - A new date/time function, `last_day()`, for finding the date corresponding to the last day of a particular month. See [Impala Date and Time Functions](#) for details.
 - New syntax, `REFRESH FUNCTIONS db_name`, lets Impala recognize newly added functions, such as UDFs created through Hive. Impala scans the metadata for a specified database to locate the new functions, which is faster and more convenient than doing a full `INVALIDATE METADATA` operation.
 - The new `TABLESAMPLE` clause for the `SELECT` statement lets you specify a percentage of the overall data volume within a table to process during a query.
 - Startup flags for the `impalad` daemon, `is_executor` and `is_coordinator`, let you divide the work on a large, busy cluster between a small number of hosts acting as query coordinators, and a larger number of hosts acting as query executors. By default, each host can act in both roles, potentially introducing bottlenecks during heavily concurrent workloads. See [How to Configure Impala with Dedicated Coordinators](#) for details.
 - A new query option, `DEFAULT_JOIN_DISTRIBUTION_MODE`, lets you change the default assumption about how join queries should handle tables with no statistics. This can help to avoid out-of-memory conditions for join queries, without manual tuning to add the `/* +SHUFFLE */` hint for queries on large tables with missing statistics.
 - The `SORT BY` clause lets you create Parquet files with more efficient compression and smaller ranges of values for specified columns, allowing Impala to apply optimizations to skip reading data from Parquet files that do not contain any values that match equality and range operators in the `WHERE` clause. See [CREATE TABLE Statement](#) for details.
 - The `max_audit_event_log_files` lets you perform log rotation for the audit event log files, similar to the rotation for regular Impala log files.
 - Kudu enhancements:
 - The `ALTER TABLE` statement can specify more attributes for a Kudu table with the `ADD COLUMNS` clause. Now you can specify `[NOT] NULL`, `ENCODING COMPRESSION`, `DEFAULT`, and `BLOCK_SIZE`. See [ALTER TABLE Statement](#) for details.
 - The `TIMESTAMP` type is now available for Kudu tables.
-  **Note:** See [Handling Date, Time, or Timestamp Data with Kudu](#) for information about the tradeoffs between performance and convenience when using this data type. For high-performance applications, you might continue to use the `BIGINT` type to represent date/time values.
- The `INSERT` and `CREATE TABLE AS SELECT` statements are more efficient when writing to Kudu tables. Formerly, the overhead for the write operations could result in timeouts when writing large numbers of rows in a single operation.

New Features in CDH 5.12 / Impala 2.9

For the full list of issues closed in this release, including the issues marked as “new features”, see the [changelog for CDH 5.12](#).

The following are some of the most significant new features in this release:

- A new function, `replace()`, which is faster than `regexp_replace()` for simple string substitutions. See [Impala String Functions](#) for details.
- Startup flags for the `impalad` daemon, `is_executor` and `is_coordinator`, let you divide the work on a large, busy cluster between a small number of hosts acting as query coordinators, and a larger number of hosts acting

as query executors. By default, each host can act in both roles, potentially introducing bottlenecks during heavily concurrent workloads. See [How to Configure Impala with Dedicated Coordinators](#) for details.

New Features in CDH 5.11.x

The following are some of the most significant new features in this release:

- Improved security for Kudu tables. This enhancement comes from upgrading the Kudu client code shipped with CDH 5.11.
 - Encryption on the wire, and in the web UI. (Not on the client side.)
 - Kerberos authentication.

New Features in Impala 2.8.x / CDH 5.10.x

- Performance and scalability improvements:
 - Improved scalability for highly concurrent loads by reducing the possibility of TCP/IP timeouts. A configuration setting, `accepted_cnxn_queue_depth`, can be adjusted upwards to avoid this type of timeout on large clusters.
 - Several performance improvements were made to the mechanism for generating native code:
 - Some queries involving analytic functions can take better advantage of native code generation.
 - Modules produced during intermediate code generation are organized to be easier to cache and reuse during the lifetime of a long-running or complicated query.
 - The `COMPUTE STATS` statement is more efficient (less time for the codegen phase) for tables with a large number of columns, especially for tables containing `TIMESTAMP` columns.
 - The logic for determining whether or not to use a runtime filter is more reliable, and the evaluation process itself is faster because of native code generation.
 - The `MT_DOP` query option enables multithreading for a number of Impala operations. `COMPUTE STATS` statements for Parquet tables use a default of `MT_DOP=4` to improve the intra-node parallelism and CPU efficiency of this data-intensive operation. See [MT_DOP Query Option](#) for details.
 - A new hint, `CLUSTERED`, allows Impala `INSERT` operations on a Parquet table that use dynamic partitioning to process a high number of partitions in a single statement. The data is ordered based on the partition key columns, and each partition is only written by a single host, reducing the amount of memory needed to buffer Parquet data while the data blocks are being constructed.
 - The new configuration setting `inc_stats_size_limit_bytes` lets you reduce the load on the catalog server when running the `COMPUTE INCREMENTAL STATS` statement for very large tables.
 - Impala folds many constant expressions within query statements, rather than evaluating them for each row. This optimization is especially useful when using functions to manipulate and format `TIMESTAMP` values, such as the result of an expression such as `to_date(now() - interval 1 day)`.
 - Parsing of complicated expressions is faster. This speedup is especially useful for queries containing large `CASE` expressions.
 - Evaluation is faster for `IN` operators with many constant arguments. The same performance improvement applies to other functions with many constant arguments.
 - Impala optimizes identical comparison operators within multiple `OR` blocks.
 - The reporting for wall-clock times and total CPU time in profile output is more accurate.
 - A new query option, `SCRATCH_LIMIT`, lets you restrict the amount of space used when a query exceeds the memory limit and activates the “spill to disk” mechanism. This option helps to avoid runaway queries or make

queries “fail fast” if they require more memory than anticipated. You can prevent runaway queries from using excessive amounts of spill space, without restarting the cluster to turn the spilling feature off entirely. See [SCRATCH_LIMIT Query Option](#) for details.

- Integration with Apache Kudu:
 - The experimental Impala support for the Kudu storage layer has been folded into the main Impala development branch. Impala can now directly access Kudu tables, opening up new capabilities such as enhanced DML operations and continuous ingestion.
 - The `DELETE` statement is a flexible way to remove data from a Kudu table. Previously, removing data from an Impala table involved removing or rewriting the underlying data files, dropping entire partitions, or rewriting the entire table. This Impala statement only works for Kudu tables.
 - The `UPDATE` statement is a flexible way to modify data within a Kudu table. Previously, updating data in an Impala table involved replacing the underlying data files, dropping entire partitions, or rewriting the entire table. This Impala statement only works for Kudu tables.
 - The `UPSERT` statement is a flexible way to ingest, modify, or both data within a Kudu table. Previously, ingesting data that might contain duplicates involved an inefficient multi-stage operation, and there was no built-in protection against duplicate data. The `UPSERT` statement, in combination with the primary key designation for Kudu tables, lets you add or replace rows in a single operation, and automatically avoids creating any duplicate data.
 - The `CREATE TABLE` statement gains some new clauses that are specific to Kudu tables: `PARTITION BY`, `PARTITIONS`, `STORED AS KUDU`, and column attributes `PRIMARY KEY`, `NULL` and `NOT NULL`, `ENCODING`, `COMPRESSION`, `DEFAULT`, and `BLOCK_SIZE`. These clauses replace the explicit `TBLPROPERTIES` settings that were required in the early experimental phases of integration between Impala and Kudu.
 - The `ALTER TABLE` statement can change certain attributes of Kudu tables. You can add, drop, or rename columns. You can add or drop range partitions. You can change the `TBLPROPERTIES` value to rename or point to a different underlying Kudu table, independently from the Impala table name in the metastore database. You cannot change the data type of an existing column in a Kudu table.
 - The `SHOW PARTITIONS` statement displays information about the distribution of data between partitions in Kudu tables. A new variation, `SHOW RANGE PARTITIONS`, displays information about the Kudu-specific partitions that apply across ranges of key values.
 - Not all Impala data types are supported in Kudu tables. In particular, currently the Impala `TIMESTAMP` type is not allowed in a Kudu table. Impala does not recognize the `UNIXTIME_MICROS` Kudu type when it is present in a Kudu table. (These two representations of date/time data use different units and are not directly compatible.) You cannot create columns of type `TIMESTAMP`, `DECIMAL`, `VARCHAR`, or `CHAR` within a Kudu table. Within a query, you can cast values in a result set to these types. Certain types, such as `BOOLEAN`, cannot be used as primary key columns.
 - Currently, Kudu tables are not interchangeable between Impala and Hive the way other kinds of Impala tables are. Although the metadata for Kudu tables is stored in the metastore database, currently Hive cannot access Kudu tables.
 - The `INSERT` statement works for Kudu tables. The organization of the Kudu data makes it more efficient than with HDFS-backed tables to insert data in small batches, such as with the `INSERT . . . VALUES` syntax.
 - Some audit data is recorded for data governance purposes. All `UPDATE`, `DELETE`, and `UPSERT` statements are characterized as `INSERT` operations in the audit log. Currently, lineage metadata is not generated for `UPDATE` and `DELETE` operations on Kudu tables.
 - Currently, Kudu tables have limited support for Sentry:
 - Access to Kudu tables must be granted to roles as usual.

- Currently, access to a Kudu table through Sentry is “all or nothing”. You cannot enforce finer-grained permissions such as at the column level, or permissions on certain operations such as `INSERT`.
- Only users with `ALL` privileges on `SERVER` can create external Kudu tables.

Because non-SQL APIs can access Kudu data without going through Sentry authorization, currently the Sentry support is considered preliminary.

- Equality and `IN` predicates in Impala queries are pushed to Kudu and evaluated efficiently by the Kudu storage layer.
- **Security:**
 - Impala can take advantage of the S3 encrypted credential store, to avoid exposing the secret key when accessing data stored on S3.
- [\[IMPALA-1654\]](#) Several kinds of DDL operations can now work on a range of partitions. The partitions can be specified using operators such as `<`, `>=`, and `!=` rather than just an equality predicate applying to a single partition. This new feature extends the syntax of several clauses of the `ALTER TABLE` statement (`DROP PARTITION`, `SET [UN]CACHED`, `SET FILEFORMAT | SERDEPROPERTIES | TBLPROPERTIES`), the `SHOW FILES` statement, and the `COMPUTE INCREMENTAL STATS` statement. It does not apply to statements that are defined to only apply to a single partition, such as `LOAD DATA`, `ALTER TABLE ... ADD PARTITION`, `SET LOCATION`, and `INSERT` with a static partitioning clause.
- The `instr()` function has optional second and third arguments, representing the character to position to begin searching for the substring, and the Nth occurrence of the substring to find.
- Improved error handling for malformed Avro data. In particular, incorrect precision or scale for `DECIMAL` types is now handled.
- Impala debug web UI:
 - In addition to “inflight” and “finished” queries, the web UI now also includes a section for “queued” queries.
 - The `/sessions` tab now clarifies how many of the displayed sections are active, and lets you sort by **Expired** status to distinguish active sessions from expired ones.
- Improved stability when DDL operations such as `CREATE DATABASE` or `DROP DATABASE` are run in Hive at the same time as an Impala `INVALIDATE METADATA` statement.
- The “out of memory” error report was made more user-friendly, with additional diagnostic information to help identify the spot where the memory limit was exceeded.
- Improved disk space usage for Java-based UDFs. Temporary copies of the associated JAR files are removed when no longer needed, so that they do not accumulate across restarts of the `catalogd` daemon and potentially cause an out-of-space condition. These temporary files are also created in the directory specified by the `local_library_dir` configuration setting, so that the storage for these temporary files can be independent from any capacity limits on the `/tmp` filesystem.

New Features in Impala 2.7.x / CDH 5.9.x

- Performance improvements:
 - [\[IMPALA-3206\]](#) Speedup for queries against `DECIMAL` columns in Avro tables. The code that parses `DECIMAL` values from Avro now uses native code generation.
 - [\[IMPALA-3674\]](#) Improved efficiency in LLVM code generation can reduce codegen time, especially for short queries.
 - [\[IMPALA-2979\]](#) Improvements to scheduling on worker nodes, enabled by the `REPLICA_PREFERENCE` query option. See [REPLICA_PREFERENCE Query Option \(CDH 5.9 or higher only\)](#) for details.

- [\[IMPALA-1683\]](#) The `REFRESH` statement can be applied to a single partition, rather than the entire table. See [REFRESH Statement](#) and [Refreshing a Single Partition](#) for details.
- Improvements to the Impala web user interface:
 - [\[IMPALA-2767\]](#) You can now force a session to expire by clicking a link in the web UI, on the `/sessions` tab.
 - [\[IMPALA-3715\]](#) The `/memz` tab includes more information about Impala memory usage.
 - [\[IMPALA-3716\]](#) The **Details** page for a query now includes a **Memory** tab.
- [\[IMPALA-3499\]](#) Scalability improvements to the catalog server. Impala handles internal communication more efficiently for tables with large numbers of columns and partitions, where the size of the metadata exceeds 2 GiB.
- [\[IMPALA-3677\]](#) You can send a `SIGUSR1` signal to any Impala-related daemon to write a Breakpad minidump. For advanced troubleshooting, you can now produce a minidump without triggering a crash. See [Breakpad Minidumps for Impala \(CDH 5.8 or higher only\)](#) for details about the Breakpad minidump feature.
- [\[IMPALA-3687\]](#) The schema reconciliation rules for Avro tables have changed slightly for `CHAR` and `VARCHAR` columns. Now, if the definition of such a column is changed in the Avro schema file, the column retains its `CHAR` or `VARCHAR` type as specified in the SQL definition, but the column name and comment from the Avro schema file take precedence. See [Creating Avro Tables](#) for details about column definitions in Avro tables.
- [\[IMPALA-3575\]](#) Some network operations now have additional timeout and retry settings. The extra configuration helps avoid failed queries for transient network problems, to avoid hangs when a sender or receiver fails in the middle of a network transmission, and to make cancellation requests more reliable despite network issues.

New Features in Impala 2.6.x / CDH 5.8.x

- Improvements to Impala support for the Amazon S3 filesystem:
 - Impala can now write to S3 tables through the `INSERT` or `LOAD DATA` statements. See [Using Impala with the Amazon S3 Filesystem](#) for general information about using Impala with S3.
 - A new query option, `S3_SKIP_INSERT_STAGING`, lets you trade off between fast `INSERT` performance and slower `INSERTS` that are more consistent if a problem occurs during the statement. The new behavior is enabled by default. See [S3_SKIP_INSERT_STAGING Query Option \(CDH 5.8 or higher only\)](#) for details about this option.
- Performance improvements for the runtime filtering feature:
 - The default for the `RUNTIME_FILTER_MODE` query option is changed to `GLOBAL` (the highest setting). See [RUNTIME_FILTER_MODE Query Option \(CDH 5.7 or higher only\)](#) for details about this option.
 - The `RUNTIME_BLOOM_FILTER_SIZE` setting is now only used as a fallback if statistics are not available; otherwise, Impala uses the statistics to estimate the appropriate size to use for each filter. See [RUNTIME_BLOOM_FILTER_SIZE Query Option \(CDH 5.7 or higher only\)](#) for details about this option.
 - New query options `RUNTIME_FILTER_MIN_SIZE` and `RUNTIME_FILTER_MAX_SIZE` let you fine-tune the sizes of the Bloom filter structures used for runtime filtering. If the filter size derived from Impala internal estimates or from the `RUNTIME_FILTER_BLOOM_SIZE` falls outside the size range specified by these options, any too-small filter size is adjusted to the minimum, and any too-large filter size is adjusted to the maximum. See [RUNTIME_FILTER_MIN_SIZE Query Option \(CDH 5.8 or higher only\)](#) and [RUNTIME_FILTER_MAX_SIZE Query Option \(CDH 5.8 or higher only\)](#) for details about these options.
 - Runtime filter propagation now applies to all the operands of `UNION` and `UNION ALL` operators.
 - Runtime filters can now be produced during join queries even when the join processing activates the spill-to-disk mechanism.

See [Runtime Filtering for Impala Queries \(CDH 5.7 or higher only\)](#) for general information about the runtime filtering feature.

- Admission control and dynamic resource pools are enabled by default. See [Admission Control and Query Queuing](#) for details about admission control.
- Impala can now manually set column statistics, using the `ALTER TABLE` statement with a `SET COLUMN STATS` clause. See [ALTER TABLE SET COLUMN STATS](#) for details.
- Impala can now write lightweight “minidump” files, rather than large core files, to save diagnostic information when any of the Impala-related daemons crash. This feature uses the open source `breakpad` framework. See [Breakpad Minidumps for Impala \(CDH 5.8 or higher only\)](#) for details.
- New query options improve interoperability with Parquet files:
 - The `PARQUET_FALLBACK_SCHEMA_RESOLUTION` query option lets Impala locate columns within Parquet files based on column name rather than ordinal position. This enhancement improves interoperability with applications that write Parquet files with a different order or subset of columns than are used in the Impala table. See [PARQUET_FALLBACK_SCHEMA_RESOLUTION Query Option \(CDH 5.8 or higher only\)](#) for details.
 - The `PARQUET_ANNOTATE_STRINGS_UTF8` query option makes Impala include the UTF-8 annotation metadata for `STRING`, `CHAR`, and `VARCHAR` columns in Parquet files created by `INSERT` or `CREATE TABLE AS SELECT` statements. See [PARQUET_ANNOTATE_STRINGS_UTF8 Query Option \(CDH 5.8 or higher only\)](#) for details.

See [Using the Parquet File Format with Impala Tables](#) for general information about working with Parquet files.

- Improvements to security and reduction in overhead for secure clusters:
 - Overall performance improvements for secure clusters. (TPC-H queries on a secure cluster were benchmarked at roughly 3x as fast as the previous release.)
 - Impala now recognizes the `auth_to_local` setting, specified through the HDFS configuration setting `hadoop.security.auth_to_local`. This feature is disabled by default; to enable it, specify `--load_auth_to_local_rules=true` in the `impalad` configuration settings. See [Mapping Kerberos Principals to Short Names for Impala](#) for details.
 - Timing improvements in the mechanism for the `impalad` daemon to acquire Kerberos tickets. This feature spreads out the overhead on the KDC during Impala startup, especially for large clusters.
 - For Kerberized clusters, the Catalog service now uses the Kerberos principal instead of the operating system user that runs the `catalogd` daemon. This eliminates the requirement to configure a `hadoop.user.group.static.mapping.overrides` setting to put the OS user into the Sentry administrative group, on clusters where the principal and the OS user name for this user are different.
- Overall performance improvements for join queries, by using a prefetching mechanism while building the in-memory hash table to evaluate join predicates. See [PREFETCH_MODE Query Option \(CDH 5.8 or higher only\)](#) for the query option to control this optimization.
- The `impala-shell` interpreter has a new command, `SOURCE`, that lets you run a set of SQL statements or other `impala-shell` commands stored in a file. You can run additional `SOURCE` commands from inside a file, to set up flexible sequences of statements for use cases such as schema setup, ETL, or reporting. See [impala-shell Command Reference](#) for details and [Running Commands and SQL Statements in impala-shell](#) for examples.
- The `millisecond()` built-in function lets you extract the fractional seconds part of a `TIMESTAMP` value. See [Impala Date and Time Functions](#) for details.
- New built-in functions `base64encode()` and `base64decode()` let you encode arbitrary string values in plaintext format. See [Impala String Functions](#) for details.
- If an Avro table is created without column definitions in the `CREATE TABLE` statement, and columns are later added through `ALTER TABLE`, the resulting table is now queryable. Missing values from the newly added columns now default to `NULL`. See [Using the Avro File Format with Impala Tables](#) for general details about working with Avro files.

- The mechanism for interpreting `DECIMAL` literals is improved, no longer going through an intermediate conversion step to `DOUBLE`:
 - Casting a `DECIMAL` value to `TIMESTAMP DOUBLE` produces a more precise value for the `TIMESTAMP` than formerly.
 - Certain function calls involving `DECIMAL` literals now succeed, when formerly they failed due to lack of a function signature with a `DOUBLE` argument.
 - Faster runtime performance for `DECIMAL` constant values, through improved native code generation for all combinations of precision and scale.

See [DECIMAL Data Type](#) for details about the `DECIMAL` type.

- Improved type accuracy for `CASE` return values. If all `WHEN` clauses of the `CASE` expression are of `CHAR` type, the final result is also `CHAR` instead of being converted to `STRING`. See [Impala Conditional Functions](#) for details about the `CASE` function.
- Uncorrelated queries using the `NOT EXISTS` operator are now supported. Formerly, the `NOT EXISTS` operator was only available for correlated subqueries.
- Improved performance for reading Parquet files.
- Improved performance for **top-N** queries, that is, those including both `ORDER BY` and `LIMIT` clauses.
- Impala optionally skips an arbitrary number of header lines from text input files on HDFS based on the `skip.header.line.count` value in the `TBLPROPERTIES` field of the table metadata. See [Data Files for Text Tables](#) for details.
- Trailing comments are now allowed in queries processed by the `impala-shell` options `-q` and `-f`.
- Impala can run `COUNT` queries for RCFile tables that include complex type columns. See [Complex Types \(CDH 5.5 or higher only\)](#) for general information about working with complex types, and [ARRAY Complex Type \(CDH 5.5 or higher only\)](#), [MAP Complex Type \(CDH 5.5 or higher only\)](#), and [STRUCT Complex Type \(CDH 5.5 or higher only\)](#) for syntax details of each type.

New Features in Impala 2.5.x / CDH 5.7.x



Note: Impala 2.5.x is available as part of CDH 5.7.x and is not available for CDH 4. Cloudera does not intend to release future versions of Impala for CDH 4 outside patch and maintenance releases if required. Given the end-of-maintenance status for CDH 4, Cloudera recommends all customers to migrate to a recent CDH 5 release.

- Dynamic partition pruning. When a query refers to a partition key column in a `WHERE` clause, and the exact set of column values are not known until the query is executed, Impala evaluates the predicate and skips the I/O for entire partitions that are not needed. For example, if a table was partitioned by year, Impala would apply this technique to a query such as `SELECT c1 FROM partitioned_table WHERE year = (SELECT MAX(year) FROM other_table)`.

The dynamic partition pruning optimization technique lets Impala avoid reading data files from partitions that are not part of the result set, even when that determination cannot be made in advance. This technique is especially valuable when performing join queries involving partitioned tables. For example, if a join query includes an `ON` clause and a `WHERE` clause that refer to the same columns, the query can find the set of column values that match the `WHERE` clause, and only scan the associated partitions when evaluating the `ON` clause.

Dynamic partition pruning is controlled by the same settings as the runtime filtering feature. By default, this feature is enabled at a medium level, because the maximum setting can use slightly more memory for queries than in previous releases. To fully enable this feature, set the query option `RUNTIME_FILTER_MODE=GLOBAL`.

- Runtime filtering. This is a wide-ranging set of optimizations that are especially valuable for join queries. Using the same technique as with dynamic partition pruning, Impala uses the predicates from `WHERE` and `ON` clauses to

determine the subset of column values from one of the joined tables could possibly be part of the result set. Impala sends a compact representation of the filter condition to the hosts in the cluster, instead of the full set of values or the entire table.

By default, this feature is enabled at a medium level, because the maximum setting can use slightly more memory for queries than in previous releases. To fully enable this feature, set the query option `RUNTIME_FILTER_MODE=GLOBAL`.

This feature involves some new query options: `RUNTIME_FILTER_MODE`, `MAX_NUM_RUNTIME_FILTERS`, `RUNTIME_BLOOM_FILTER_SIZE`, `RUNTIME_FILTER_WAIT_TIME_MS`, and `DISABLE_ROW_RUNTIME_FILTERING`.

- More efficient use of the HDFS caching feature, to avoid hotspots and bottlenecks that could occur if heavily used cached data blocks were always processed by the same host. By default, Impala now randomizes which host processes each cached HDFS data block, when cached replicas are available on multiple hosts. (Remember to use the `WITH REPLICATION` clause with the `CREATE TABLE` or `ALTER TABLE` statement when enabling HDFS caching for a table or partition, to cache the same data blocks across multiple hosts.) The new query option `SCHEDULE_RANDOM_REPLICA` lets you fine-tune the interaction with HDFS caching even more.
- The `TRUNCATE TABLE` statement now accepts an `IF EXISTS` clause, making `TRUNCATE TABLE` easier to use in setup or ETL scripts where the table might or might not exist.
- Improved performance and reliability for the `DECIMAL` data type:
 - Using `DECIMAL` values in a `GROUP BY` clause now triggers the native code generation optimization, speeding up queries that group by values such as prices.
 - Checking for overflow in `DECIMAL` multiplication is now substantially faster, making `DECIMAL` a more practical data type in some use cases where formerly `DECIMAL` was much slower than `FLOAT` or `DOUBLE`.
 - Multiplying a mixture of `DECIMAL` and `FLOAT` or `DOUBLE` values now returns the `DOUBLE` rather than `DECIMAL`. This change avoids some cases where an intermediate value would underflow or overflow and become `NULL` unexpectedly.
- For UDFs written in Java, or Hive UDFs reused for Impala, Impala now allows parameters and return values to be primitive types. Formerly, these things were required to be one of the “Writable” object types.
- Performance improvements for HDFS I/O. Impala now caches HDFS file handles to avoid the overhead of repeatedly opening the same file.
- Performance improvements for queries involving nested complex types. Certain basic query types, such as counting the elements of a complex column, now use an optimized code path.
- Improvements to the memory reservation mechanism for the Impala admission control feature. You can specify more settings, such as the timeout period and maximum aggregate memory used, for each resource pool instead of globally for the Impala instance. The default limit for concurrent queries (the `max requests` setting) is now unlimited instead of 200. The Cloudera Manager user interface for admission control has been reworked, with the settings available under the **Dynamic Resource Pools** window.
- Performance improvements related to code generation. Even in queries where code generation is not performed for some phases of execution (such as reading data from Parquet tables), Impala can still use code generation in other parts of the query, such as evaluating functions in the `WHERE` clause.
- Performance improvements for queries using aggregation functions on high-cardinality columns. Formerly, Impala could do unnecessary extra work to produce intermediate results for operations such as `DISTINCT` or `GROUP BY` on columns that were unique or had few duplicate values. Now, Impala decides at run time whether it is more efficient to do an initial aggregation phase and pass along a smaller set of intermediate data, or to pass raw intermediate data back to next phase of query processing to be aggregated there. This feature is known as **streaming pre-aggregation**. In case of performance regression, this feature can be turned off using the `DISABLE_STREAMING_PREAGGREGATIONS` query option.
- Spill-to-disk feature now always recommended. In earlier releases, the spill-to-disk feature could be turned off using a pair of configuration settings, `enable_partitioned_aggregation=false` and

`enable_partitioned_hash_join=false`. The latest improvements in the spill-to-disk mechanism, and related features that interact with it, make this feature robust enough that disabling it is now no longer needed or supported. In particular, some new features in CDH 5.7 / Impala 2.5 and higher do not work when the spill-to-disk feature is disabled.

- Improvements to scripting capability for the `impala-shell` command, through user-specified substitution variables that can appear in statements processed by `impala-shell`:
 - The `--var` command-line option lets you pass key-value pairs to `impala-shell`. The shell can substitute the values into queries before executing them, where the query text contains the notation `${var:varname}`. For example, you might prepare a SQL file containing a set of DDL statements and queries containing variables for database and table names, and then pass the applicable names as part of the `impala-shell -f filename` command.
 - The `SET` and `UNSET` commands within the `impala-shell` interpreter now work with user-specified substitution variables, as well as the built-in query options. The two kinds of variables are divided in the `SET` output. As with variables defined by the `--var` command-line option, you refer to the user-specified substitution variables in queries by using the notation `${var:varname}` in the query text. Because the substitution variables are processed by `impala-shell` instead of the `impalad` backend, you cannot define your own substitution variables through the `SET` statement in a JDBC or ODBC application.
- Performance improvements for query startup. Impala better parallelizes certain work when coordinating plan distribution between `impalad` instances, which improves startup time for queries involving tables with many partitions on large clusters, or complicated queries with many plan fragments.
- Performance and scalability improvements for tables with many partitions. The memory requirements on the coordinator node are reduced, making it substantially faster and less resource-intensive to do joins involving several tables with thousands of partitions each.
- Whitelisting for access to internal APIs. For applications that need direct access to Impala APIs, without going through the HiveServer2 or Beeswax interfaces, you can specify a list of Kerberos users who are allowed to call those APIs. By default, the `impala` and `hdfs` users are the only ones authorized for this kind of access. Any users not explicitly authorized through the `internal_principals_whitelist` configuration setting are blocked from accessing the APIs. This setting applies to all the Impala-related daemons, although currently it is primarily used for HDFS to control the behavior of the catalog server.
- Improvements to Impala integration and usability for Hue. (The code changes are actually on the Hue side.)
 - The list of tables now refreshes dynamically.
- Usability improvements for case-insensitive queries. You can now use the operators `ILIKE` and `IREGEXP` to perform case-insensitive wildcard matches or regular expression matches, rather than explicitly converting column values with `UPPER` or `LOWER`.
- Performance and reliability improvements for DDL and insert operations on partitioned tables with a large number of partitions. Impala only re-evaluates metadata for partitions that are affected by a DDL operation, not all partitions in the table. While a DDL or insert statement is in progress, other Impala statements that attempt to modify metadata for the same table wait until the first one finishes.
- Reliability improvements for the `LOAD DATA` statement. Previously, this statement would fail if the source HDFS directory contained any subdirectories at all. Now, the statement ignores any hidden subdirectories, for example `_impala_insert_staging`.
- A new operator, `IS [NOT] DISTINCT FROM`, lets you compare values and always get a `true` or `false` result, even if one or both of the values are `NULL`. The `IS NOT DISTINCT FROM` operator, or its equivalent `<=>` notation, improves the efficiency of join queries that treat key values that are `NULL` in both tables as equal.
- Security enhancements for the `impala-shell` command. A new option, `--ldap_password_cmd`, lets you specify a command to retrieve the LDAP password. The resulting password is then used to authenticate the `impala-shell` command with the LDAP server.

- The `CREATE TABLE AS SELECT` statement now accepts a `PARTITIONED BY` clause, which lets you create a partitioned table and insert data into it with a single statement.
- User-defined functions (UDFs and UDAFs) written in C++ now persist automatically when the `catalogd` daemon is restarted. You no longer have to run the `CREATE FUNCTION` statements again after a restart.
- User-defined functions (UDFs) written in Java can now persist when the `catalogd` daemon is restarted, and can be shared transparently between Impala and Hive. You must do a one-time operation to recreate these UDFs using new `CREATE FUNCTION` syntax, without a signature for arguments or the return value. Afterwards, you no longer have to run the `CREATE FUNCTION` statements again after a restart. Although Impala does not have visibility into the UDFs that implement the Hive built-in functions, user-created Hive UDFs are now automatically available for calling through Impala.
- Reliability enhancements for memory management. Some aggregation and join queries that formerly might have failed with an out-of-memory error due to memory contention, now can succeed using the spill-to-disk mechanism.
- The `SHOW DATABASES` statement now returns two columns rather than one. The second column includes the associated comment string, if any, for each database. Adjust any application code that examines the list of databases and assumes the result set contains only a single column.
- A new optimization speeds up aggregation operations that involve only the partition key columns of partitioned tables. For example, a query such as `SELECT COUNT(DISTINCT k), MIN(k), MAX(k) FROM t1` can avoid reading any data files if `T1` is a partitioned table and `K` is one of the partition key columns. Because this technique can produce different results in cases where HDFS files in a partition are manually deleted or are empty, you must enable the optimization by setting the query option `OPTIMIZE_PARTITION_KEY_SCANS`.
- The `DESCRIBE` statement can now display metadata about a database, using the syntax `DESCRIBE DATABASE db_name`.
- The `uuid()` built-in function generates an alphanumeric value that you can use as a guaranteed unique identifier. The uniqueness applies even across tables, for cases where an ascending numeric sequence is not suitable.

New Features in Impala 2.4.x / CDH 5.6.x



Note: Impala 2.4.0 is available as part of CDH 5.6.0 and is not available for CDH 4. Cloudera does not intend to release future versions of Impala for CDH 4 outside patch and maintenance releases if required. Given the end-of-maintenance status for CDH 4, Cloudera recommends all customers to migrate to a recent CDH 5 release.

- From a user perspective, the Impala features are the same as in CDH 5.5 / Impala 2.3.

New Features in Impala 2.3.x / CDH 5.5.x



Note: Impala 2.3.x is available as part of CDH 5.5.x and is not available for CDH 4. Cloudera does not intend to release future versions of Impala for CDH 4 outside patch and maintenance releases if required. Given the end-of-maintenance status for CDH 4, Cloudera recommends all customers to migrate to a recent CDH 5 release.

The following are the major new features in Impala 2.3.x. This major release, available as part of CDH 5.5.x, contains improvements to SQL syntax (particularly new support for complex types), performance, manageability, security.

- Complex data types: `STRUCT`, `ARRAY`, and `MAP`. These types can encode multiple named fields, positional items, or key-value pairs within a single column. You can combine these types to produce nested types with arbitrarily deep nesting, such as an `ARRAY` of `STRUCT` values, a `MAP` where each key-value pair is an `ARRAY` of other `MAP` values, and so on. Currently, complex data types are only supported for the Parquet file format.
- Column-level authorization lets you define access to particular columns within a table, rather than the entire table. This feature lets you reduce the reliance on creating views to set up authorization schemes for subsets of

information. See [Hive SQL Syntax for Use with Sentry](#) for background details, and [GRANT Statement \(CDH 5.2 or higher only\)](#) and [REVOKE Statement \(CDH 5.2 or higher only\)](#) for Impala-specific syntax.

- The `TRUNCATE TABLE` statement removes all the data from a table without removing the table itself.
- Nested loop join queries. Some join queries that formerly required equality comparisons can now use operators such as `<` or `>=`. This same join mechanism is used internally to optimize queries that retrieve values from complex type columns.
- Reduced memory usage and improved performance and robustness for spill-to-disk feature.
- Performance improvements for querying Parquet data files containing multiple row groups and multiple data blocks:
 - For files written by Hive, SparkSQL, and other Parquet MR writers and spanning multiple HDFS blocks, Impala now scans the extra data blocks locally when possible, rather than using remote reads.
 - Impala queries benefit from the improved alignment of row groups with HDFS blocks for Parquet files written by Hive, MapReduce, and other components in CDH 5.5 and higher. (Impala itself never writes multiblock Parquet files, so the alignment change does not apply to Parquet files produced by Impala.) These Parquet writers now add padding to Parquet files that they write to align row groups with HDFS blocks. The `parquet.writer.max-padding` setting specifies the maximum number of bytes, by default 8 megabytes, that can be added to the file between row groups to fill the gap at the end of one block so that the next row group starts at the beginning of the next block. If the gap is larger than this size, the writer attempts to fit another entire row group in the remaining space. Include this setting in the `hive-site` configuration file to influence Parquet files written by Hive, or the `hdfs-site` configuration file to influence Parquet files written by all non-Impala components.
- Many new built-in scalar functions, for convenience and enhanced portability of SQL that uses common industry extensions.

Math functions:

- `ATAN2`
- `COSH`
- `COT`
- `DCEIL`
- `DEXP`
- `DFLOOR`
- `DLOG10`
- `DPOW`
- `DROUND`
- `DSQRT`
- `DTRUNC`
- `FACTORIAL`, and corresponding `!` operator
- `FPOW`
- `RADIANS`
- `RANDOM`
- `SINH`
- `TANH`

String functions:

- `BTRIM`
- `CHR`
- `REGEXP_LIKE`
- `SPLIT_PART`

Date and time functions:

- INT_MONTHS_BETWEEN
- MONTHS_BETWEEN
- TIMEOFDAY
- TIMESTAMP_CMP

Bit manipulation functions:

- BITAND
- BITNOT
- BITOR
- BITXOR
- COUNTSET
- GETBIT
- ROTATELEFT
- ROTATERIGHT
- SETBIT
- SHIFLEFT
- SHIFTRIGHT

Type conversion functions:

- TYPEOF

The `effective_user()` function.

- New built-in analytic functions: `PERCENT_RANK`, `NTILE`, `CUME_DIST`.
- The `DROP DATABASE` statement now works for a non-empty database. When you specify the optional `CASCADE` clause, any tables in the database are dropped before the database itself is removed.
- The `DROP TABLE` and `ALTER TABLE DROP PARTITION` statements have a new optional keyword, `PURGE`. This keyword causes Impala to immediately remove the relevant HDFS data files rather than sending them to the HDFS trashcan. This feature can help to avoid out-of-space errors on storage devices, and to avoid files being left behind in case of a problem with the HDFS trashcan, such as the trashcan not being configured or being in a different HDFS encryption zone than the data files.
- The `impala-shell` command has a new feature for live progress reporting. This feature is enabled through the `--live_progress` and `--live_summary` command-line options, or during a session through the `LIVE_SUMMARY` and `LIVE_PROGRESS` query options.
- The `impala-shell` command also now displays a random “tip of the day” when it starts.
- The `impala-shell` option `-f` now recognizes a special filename `-` to accept input from stdin.
- Format strings for the `unix_timestamp()` function can now include numeric timezone offsets.
- Impala can now run a specified command to obtain the password to decrypt a private-key PEM file, rather than having the private-key file be unencrypted on disk.
- Impala components now can use SSL for more of their internal communication. SSL is used for communication between all three Impala-related daemons when the configuration option `ssl_server_certificate` is enabled. SSL is used for communication with client applications when the configuration option `ssl_client_ca_certificate` is enabled.

Currently, you can only use one of server-to-server TLS/SSL encryption or Kerberos authentication. This limitation is tracked by the issue [IMPALA-2598](#).

- Improved flexibility for intermediate data types in user-defined aggregate functions (UDAFs).

In CDH 5.5.2 / Impala 2.3.2, the bug fix for [IMPALA-2598](#) removes the restriction on using both Kerberos and SSL for internal communication between Impala components.

New Features in Impala 2.2.x for CDH 5.4.3 and 5.4.4

No new features added to the Impala code. The certification of Impala with EMC Isilon under CDH 5.4.4 means that now you can query data stored on Isilon storage devices through Impala. See [Installing Cloudera Manager and CDH with Isilon](#) for details. The same level of Impala is included with both CDH 5.4.3 and 5.4.4.



Note: The Impala 2.2.x maintenance releases now use the CDH 5.4.x numbering system rather than increasing the Impala version numbers. Impala 2.2 and higher are not available under CDH 4.

New Features in Impala 2.2.x / CDH 5.4.x



Note: Impala 2.2.0 is available as part of CDH 5.4.0 and is not available for CDH 4. Cloudera does not intend to release future versions of Impala for CDH 4 outside patch and maintenance releases if required. Given the end-of-maintenance status for CDH 4, Cloudera recommends all customers to migrate to a recent CDH 5 release.

The following are the major new features in Impala 2.2.x. This release, available as part of CDH 5.4.x, contains improvements to performance, manageability, security, and SQL syntax.

- Several improvements to date and time features enable higher interoperability with Hive and other database systems, provide more flexibility for handling time zones, and future-proof the handling of `TIMESTAMP` values:
 - The `WITH REPLICATION` clause for the `CREATE TABLE` and `ALTER TABLE` statements lets you control the replication factor for HDFS caching for a specific table or partition. By default, each cached block is only present on a single host, which can lead to CPU contention if the same host processes each cached block. Increasing the replication factor lets Impala choose different hosts to process different cached blocks, to better distribute the CPU load.
 - Startup flags for the `impalad` daemon enable a higher level of compatibility with `TIMESTAMP` values written by Hive, and more flexibility for working with date and time data using the local time zone instead of UTC. To enable these features, set the `impalad` startup flags `-use_local_tz_for_unix_timestamp_conversions=true` and `-convert_legacy_hive_parquet_utc_timestamps=true`.

The `-use_local_tz_for_unix_timestamp_conversions` setting controls how the `unix_timestamp()`, `from_unixtime()`, and `now()` functions handle time zones. By default (when this setting is turned off), Impala considers all `TIMESTAMP` values to be in the UTC time zone when converting to or from Unix time values. When this setting is enabled, Impala treats `TIMESTAMP` values passed to or returned from these functions to be in the local time zone. When this setting is enabled, take particular care that all hosts in the cluster have the same timezone settings, to avoid inconsistent results depending on which host reads or writes `TIMESTAMP` data.

The `-convert_legacy_hive_parquet_utc_timestamps` setting causes Impala to convert `TIMESTAMP` values to the local time zone when it reads them from Parquet files written by Hive. This setting only applies to data using the Parquet file format, where Impala can use metadata in the files to reliably determine that the files were written by Hive. If in the future Hive changes the way it writes `TIMESTAMP` data in Parquet, Impala will automatically handle that new `TIMESTAMP` encoding.

See [TIMESTAMP Data Type](#) for details about time zone handling and the configuration options for Impala / Hive compatibility with Parquet format.

- In Impala 2.2.0 and higher, built-in functions that accept or return integers representing `TIMESTAMP` values use the `BIGINT` type for parameters and return values, rather than `INT`. This change lets the date and time functions avoid an overflow error that would otherwise occur on January 19th, 2038 (known as the [“Year 2038 problem”](#) or [“Y2K38 problem”](#)). This change affects the `FROM_UNIXTIME()` and `UNIX_TIMESTAMP()`

functions. You might need to change application code that interacts with these functions, change the types of columns that store the return values, or add `CAST()` calls to SQL statements that call these functions.

See [Impala Date and Time Functions](#) for the current function signatures.

- The `SHOW FILES` statement lets you view the names and sizes of the files that make up an entire table or a specific partition. See [SHOW FILES Statement](#) for details.
- Impala can now run queries against Parquet data containing columns with complex or nested types, as long as the query only refers to columns with scalar types.
- Performance improvements for queries that include `IN()` operators and involve partitioned tables.
- The new `-max_log_files` configuration option specifies how many log files to keep at each severity level. The default value is 10, meaning that Impala preserves the latest 10 log files for each severity level (`INFO`, `WARNING`, and `ERROR`) for each Impala-related daemon (`impalad`, `statedored`, and `catalogd`). Impala checks to see if any old logs need to be removed based on the interval specified in the `logbufsecs` setting, every 5 seconds by default. See [Rotating Impala Logs](#) for details.
- Redaction of sensitive data from Impala log files. This feature protects details such as credit card numbers or tax IDs from administrators who see the text of SQL statements in the course of monitoring and troubleshooting a Hadoop cluster. See [Redacting Sensitive Information from Impala Log Files](#) for background information for Impala users, and [How to Enable Sensitive Data Redaction](#) for usage details.
- Lineage information is available for data created or queried by Impala. This feature lets you track who has accessed data through Impala SQL statements, down to the level of specific columns, and how data has been propagated between tables. See [Viewing Lineage Information for Impala Data](#) for background information for Impala users, [Managing Hive and Impala Lineage Properties](#) for usage details, and [Using the Lineage View](#) for how to interpret the lineage information.
- Impala tables and partitions can now be located on the Amazon Simple Storage Service (S3) filesystem, for convenience in cases where data is already located in S3 and you prefer to query it in-place. Queries might have lower performance than when the data files reside on HDFS, because Impala uses some HDFS-specific optimizations. Impala can query data in S3, but cannot write to S3. Therefore, statements such as `INSERT` and `LOAD DATA` are not available when the destination table or partition is in S3. See [Using Impala with the Amazon S3 Filesystem](#) for details.



Important:

Impala query support for Amazon S3 is included in CDH 5.4 / Impala 2.2, but is not supported or recommended for production use in this version.

- Improved support for HDFS encryption. The `LOAD DATA` statement now works when the source directory and destination table are in different encryption zones.
- Additional arithmetic function `mod()`. See [Impala Mathematical Functions](#) for details.
- Flexibility to interpret `TIMESTAMP` values using the UTC time zone (the traditional Impala behavior) or using the local time zone (for compatibility with `TIMESTAMP` values produced by Hive).
- Enhanced support for ETL using tools such as Flume. Impala ignores temporary files typically produced by these tools (filenames with suffixes `.copying` and `.tmp`).
- The CPU requirement for Impala, which had become more restrictive in Impala 2.0.x and 2.1.x, has now been relaxed.

The prerequisite for CPU architecture has been relaxed in Impala 2.2.0 and higher. From this release onward, Impala works on CPUs that have the SSSE3 instruction set. The SSE4 instruction set is no longer required. This relaxed requirement simplifies the upgrade planning from Impala 1.x releases, which also worked on SSSE3-enabled processors.

- Enhanced support for `CHAR` and `VARCHAR` types in the `COMPUTE STATS` statement.
- The amount of memory required during setup for “spill to disk” operations is greatly reduced. This enhancement reduces the chance of a memory-intensive join or aggregation query failing with an out-of-memory error.
- Several new conditional functions provide enhanced compatibility when porting code that uses industry extensions. The new functions are: `isfalse()`, `isnotfalse()`, `isnottrue()`, `istrue()`, `nonnullvalue()`, and `nullvalue()`. See [Impala Conditional Functions](#) for details.
- The Impala debug web UI now can display a visual representation of the query plan. On the `/queries` tab, select **Details** for a particular query. The **Details** page includes a **Plan** tab with a plan diagram that you can zoom in or out (using scroll gestures through mouse wheel or trackpad).

New Features in Impala 2.1.8 / CDH 5.3.10

This point release is exclusively a bug fix release.



Note: This Impala maintenance release is only available as part of CDH 5, not under CDH 4.

New Features in Impala 2.1.7 / CDH 5.3.9

This point release is exclusively a bug fix release.



Note: This Impala maintenance release is only available as part of CDH 5, not under CDH 4.

New Features in Impala 2.1.6 / CDH 5.3.8

This point release is exclusively a bug fix release.



Note: This Impala maintenance release is only available as part of CDH 5, not under CDH 4.

New Features in Impala 2.1.5 / CDH 5.3.6

This point release is exclusively a bug fix release.



Note: This Impala maintenance release is only available as part of CDH 5, not under CDH 4.

New Features in Impala 2.1.0 / CDH 5.3.0

This release contains the following enhancements to query performance and system scalability:

- Impala can now collect statistics for individual partitions in a partitioned table, rather than processing the entire table for each `COMPUTE STATS` statement. This feature is known as incremental statistics, and is controlled by the `COMPUTE INCREMENTAL STATS` syntax. (You can still use the original `COMPUTE STATS` statement for nonpartitioned tables or partitioned tables that are unchanging or whose contents are entirely replaced all at once.) See [COMPUTE STATS Statement](#) and [Table and Column Statistics](#) for details.
- Optimization for small queries lets Impala process queries that process very few rows without the unnecessary overhead of parallelizing and generating native code. Reducing this overhead lets Impala clear small queries quickly, keeping YARN resources and admission control slots available for data-intensive queries. The number of rows considered to be a “small” query is controlled by the `EXEC_SINGLE_NODE_ROWS_THRESHOLD` query option. See [EXEC_SINGLE_NODE_ROWS_THRESHOLD Query Option \(CDH 5.3 or higher only\)](#) for details.

- An enhancement to the statestore component lets it transmit heartbeat information independently of broadcasting metadata updates. This optimization improves reliability of health checking on large clusters with many tables and partitions.
- The memory requirement for querying gzip-compressed text is reduced. Now Impala decompresses the data as it is read, rather than reading the entire gzipped file and decompressing it in memory.

New Features in Impala 2.0.0 / CDH 5.2.0

The following are the major new features in Impala 2.0. This major release, available both with CDH 5.2 and for CDH 4, contains improvements to performance, scalability, security, and SQL syntax.

- Queries with joins or aggregation functions involving high volumes of data can now use temporary work areas on disk, reducing the chance of failure due to out-of-memory errors. When the required memory for the intermediate result set exceeds the amount available on a particular node, the query automatically uses a temporary work area on disk. This “spill to disk” mechanism is similar to the `ORDER BY` improvement from Impala 1.4. For details, see [SQL Operations that Spill to Disk](#).
- Subquery enhancements:
 - Subqueries are now allowed in the `WHERE` clause, for example with the `IN` operator.
 - The `EXISTS` and `NOT EXISTS` operators are available. They are always used in conjunction with subqueries.
 - The `IN` and `NOT IN` queries can now operate on the result set from a subquery, not just a hardcoded list of values.
 - Uncorrelated subqueries let you compare against one or more values for equality, `IN`, and `EXISTS` comparisons. For example, you might use `WHERE` clauses such as `WHERE column = (SELECT MAX(some_other_column) FROM table)` or `WHERE column IN (SELECT some_other_column FROM table WHERE conditions)`.
 - Correlated subqueries let you cross-reference values from the outer query block and the subquery.
 - Scalar subqueries let you substitute the result of single-value aggregate functions such as `MAX()`, `MIN()`, `COUNT()`, or `AVG()`, where you would normally use a numeric value in a `WHERE` clause.

For details about subqueries, see [Subqueries in Impala SELECT Statements](#). For information about new and improved operators, see [EXISTS Operator](#) and [IN Operator](#).

- Analytic functions such as `RANK()`, `LAG()`, `LEAD()`, and `FIRST_VALUE()` let you analyze sequences of rows with flexible ordering and grouping. Existing aggregate functions such as `MAX()`, `SUM()`, and `COUNT()` can also be used in an analytic context. See [Impala Analytic Functions](#) for details. See [Impala Aggregate Functions](#) for enhancements to existing aggregate functions.
- New data types provide greater compatibility with source code from traditional database systems:
 - `VARCHAR` is like the `STRING` data type, but with a maximum length. See [VARCHAR Data Type \(CDH 5.2 or higher only\)](#) for details.
 - `CHAR` is like the `STRING` data type, but with a precise length. Short values are padded with spaces on the right. See [CHAR Data Type \(CDH 5.2 or higher only\)](#) for details.
- Security enhancements:
 - Formerly, Impala was restricted to using either Kerberos or LDAP / Active Directory authentication within a cluster. Now, Impala can freely accept either kind of authentication request, allowing you to set up some hosts with Kerberos authentication and others with LDAP or Active Directory. See [Using Multiple Authentication Methods with Impala](#) for details.
 - `GRANT` statement. See [GRANT Statement \(CDH 5.2 or higher only\)](#) for details.
 - `REVOKE` statement. See [REVOKE Statement \(CDH 5.2 or higher only\)](#) for details.
 - `CREATE ROLE` statement. See [CREATE ROLE Statement \(CDH 5.2 or higher only\)](#) for details.
 - `DROP ROLE` statement. See [DROP ROLE Statement \(CDH 5.2 or higher only\)](#) for details.
 - `SHOW ROLES` and `SHOW ROLE GRANT` statements. See [SHOW Statement](#) for details.
 - To complement the HDFS encryption feature, a new Impala configuration option, `--disk_spill_encryption` secures sensitive data from being observed or tampered with when temporarily stored on disk.

The new security-related SQL statements work along with the Sentry authorization framework. See [Enabling Sentry Authorization for Impala](#) for details.

- Impala can now read compressed text files compressed by gzip, bzip, or Snappy. These files do not require any special table settings to work in an Impala text table. Impala recognizes the compression type automatically based on file extensions of `.gz`, `.bz2`, and `.snappy` respectively. These types of compressed text files are intended for convenience with existing ETL pipelines. Their non-splittable nature means they are not optimal for high-performance parallel queries. See [Using gzip, bzip2, or Snappy-Compressed Text Files](#) for details.
- Query hints can now use comment notation, `/* +hint_name */` or `-- +hint_name`, at the same places in the query where the hints enclosed by `[]` are recognized. This enhancement makes it easier to reuse Impala queries on other database systems. See [Optimizer Hints in Impala](#) for details.
- A new query option, `QUERY_TIMEOUT_S`, lets you specify a timeout period in seconds for individual queries.

The working of the `--idle_query_timeout` configuration option is extended. If no `QUERY_OPTION_S` query option is in effect, `--idle_query_timeout` works the same as before, setting the timeout interval. When the `QUERY_OPTION_S` query option is specified, its maximum value is capped by the value of the `--idle_query_timeout` option.

That is, the system administrator sets the default and maximum timeout through the `--idle_query_timeout` startup option, and then individual users or applications can set a lower timeout value if desired through the `QUERY_TIMEOUT_S` query option. See [Setting Timeout Periods for Daemons, Queries, and Sessions](#) and [QUERY_TIMEOUT_S Query Option \(CDH 5.2 or higher only\)](#) for details.

- New functions `VAR_SAMP()` and `VAR_POP()` are aliases for the existing `VARIANCE_SAMP()` and `VARIANCE_POP()` functions.
- A new date and time function, `DATE_PART()`, provides similar functionality to `EXTRACT()`. You can also call the `EXTRACT()` function using the SQL-99 syntax, `EXTRACT(unit FROM timestamp)`. These enhancements simplify the porting process for date-related code from other systems. See [Impala Date and Time Functions](#) for details.
- New approximation features provide a fast way to get results when absolute precision is not required:
 - The `APPX_COUNT_DISTINCT` query option lets Impala rewrite `COUNT(DISTINCT)` calls to use `NDV()` instead, which speeds up the operation and allows multiple `COUNT(DISTINCT)` operations in a single query. See [APPX_COUNT_DISTINCT Query Option \(CDH 5.2 or higher only\)](#) for details.

The `APPX_MEDIAN()` aggregate function produces an estimate for the median value of a column by using sampling. See [APPX_MEDIAN Function](#) for details.

- Impala now supports a `DECODE()` function. This function works as a shorthand for a `CASE()` expression, and improves compatibility with SQL code containing vendor extensions. See [Impala Conditional Functions](#) for details.
- The `STDDEV()`, `STDDEV_POP()`, `STDDEV_SAMP()`, `VARIANCE()`, `VARIANCE_POP()`, `VARIANCE_SAMP()`, and `NDV()` aggregate functions now all return `DOUBLE` results rather than `STRING`. Formerly, you were required to `CAST()` the result to a numeric type before using it in arithmetic operations.
- The default settings for Parquet block size, and the associated `PARQUET_FILE_SIZE` query option, are changed. Now, Impala writes Parquet files with a size of 256 MB and an HDFS block size of 256 MB. Previously, Impala attempted to write Parquet files with a size of 1 GB and an HDFS block size of 1 GB. In practice, Impala used a conservative estimate of the disk space needed for each Parquet block, leading to files that were typically 512 MB anyway. Thus, this change will make the file size more accurate if you specify a value for the `PARQUET_FILE_SIZE` query option. It also reduces the amount of memory reserved during `INSERT` into Parquet tables, potentially avoiding out-of-memory errors and improving scalability when inserting data into Parquet tables.
- Anti-joins are now supported, expressed using the `LEFT ANTI JOIN` and `RIGHT ANTI JOIN` clauses. These clauses returns results from one table that have no match in the other table. You might use this type of join in the same sorts of use cases as the `NOT EXISTS` and `NOT IN` operators. See [Joins in Impala SELECT Statements](#) for details.

- The `SET` command in `impala-shell` has been promoted to a real SQL statement. You can now set query options such as `PARQUET_FILE_SIZE`, `MEM_LIMIT`, and `SYNC_DDL` within JDBC, ODBC, or any other kind of application that submits SQL without going through the `impala-shell` interpreter. See [SET Statement](#) for details.
- The `impala-shell` interpreter now reads settings from an optional configuration file, named `$HOME/.impalarc` by default. See [impala-shell Configuration File](#) for details.
- The library used for regular expression parsing has changed from Boost to Google RE2. This implementation change adds support for non-greedy matches using the `. * ?` notation. This and other changes in the way regular expressions are interpreted means you might need to re-test queries that use functions such as `regexp_extract()` or `regexp_replace()`, or operators such as `REGEXP` or `RLIKE`. See [Apache Impala Incompatible Changes and Limitations](#) on page 1925 for those details.

New Features in Impala 1.4.0 / CDH 5.1.0

The following are the major new features in Impala 1.4.

- The `DECIMAL` data type lets you store fixed-precision values, for working with currency or other fractional values where it is important to represent values exactly and avoid rounding errors. This feature includes enhancements to built-in functions, numeric literals, and arithmetic expressions.
- On CDH 5, Impala can take advantage of the HDFS caching feature to “pin” entire tables or individual partitions in memory, to speed up queries on frequently accessed data and reduce the CPU overhead of memory-to-memory copying. When HDFS files are cached in memory, Impala can read the cached data without any disk reads, and without making an additional copy of the data in memory. Other Hadoop components that read the same data files also experience a performance benefit.
- Impala can now use Sentry-based authorization based either on the original policy file, or on rules defined by `GRANT` and `REVOKE` statements issued through Hive. See [Enabling Sentry Authorization for Impala](#) for details.
- For interoperability with Parquet files created through other Hadoop components, such as Pig or MapReduce jobs, you can create an Impala table that automatically sets up the column definitions based on the layout of an existing Parquet data file.
- `ORDER BY` queries no longer require a `LIMIT` clause. If the size of the result set to be sorted exceeds the memory available to Impala, Impala uses a temporary work space on disk to perform the sort operation.
- LDAP connections can be secured through either SSL or TLS.
- The following new built-in scalar and aggregate functions are available:
 - A new built-in function, `EXTRACT()`, returns one date or time field from a `TIMESTAMP` value.
 - A new built-in function, `TRUNC()`, truncates date/time values to a particular granularity, such as year, month, day, hour, and so on.
 - `ADD_MONTHS()` built-in function, an alias for the existing `MONTHS_ADD()` function.
 - A new built-in function, `ROUND()`, rounds `DECIMAL` values to a specified number of fractional digits.
 - Several built-in aggregate functions for computing properties for statistical distributions: `STDDEV()`, `STDDEV_SAMP()`, `STDDEV_POP()`, `VARIANCE()`, `VARIANCE_SAMP()`, and `VARIANCE_POP()`.
 - Several new built-in functions, such as `MAX_INT()`, `MIN_SMALLINT()`, and so on, let you conveniently check whether data values are in an expected range. You might be able to switch a column to a smaller type, saving memory during processing.
 - New built-in functions, `IS_INF()` and `IS_NAN()`, check for the special values infinity and “not a number”. These values could be specified as `inf` or `nan` in text data files, or be produced by certain arithmetic expressions.
- The `SHOW PARTITIONS` statement displays information about the structure of a partitioned table.

- New configuration options for the `impalad` daemon let you specify initial memory usage for all queries. The initial resource requests handled by Llama and YARN can be expanded later if needed, avoiding unnecessary over-allocation and reducing the chance of out-of-memory conditions.
- The Impala `CREATE TABLE` statement now has a `STORED AS AVRO` clause, allowing you to create Avro tables through Impala.
- New `impalad` configuration options let you fine-tune the calculations Impala makes to estimate resource requirements for each query. These options can help avoid problems due to overconsumption due to too-low estimates, or underutilization due to too-high estimates.
- A new `SUMMARY` command in the `impala-shell` interpreter provides a high-level summary of the work performed at each stage of the explain plan. The summary is also included in output from the `PROFILE` command.
- Performance improvements for the `COMPUTE STATS` statement:
 - The `NDV` function is speeded up through native code generation.
 - Because the `NULL` count is not currently used by the Impala query planner, in Impala 1.4.0 and higher, `COMPUTE STATS` does not count the `NULL` values for each column. (The `#Nulls` field of the stats table is left as `-1`, signifying that the value is unknown.)
- Performance improvements for partition pruning. This feature reduces the time spent in query planning, for partitioned tables with thousands of partitions. Previously, Impala typically queried tables with up to approximately 3000 partitions. With the performance improvement in partition pruning, now Impala can comfortably handle tables with tens of thousands of partitions.
- The documentation provides additional guidance for planning tasks.
- The `impala-shell` interpreter now supports UTF-8 characters for input and output. You can control whether `impala-shell` ignores invalid Unicode code points through the `--strict_unicode` option. (Although this option is removed in Impala 2.0.)

New Features in Impala 1.3.2 / CDH 5.0.4

No new features. This point release is exclusively a bug fix release for the IMPALA-1019 issue related to HDFS caching.



Note: Impala 1.3.2 is only available as part of CDH 5.0.4, not under CDH 4.

New Features in Impala 1.3.1 / CDH 5.0.3

This point release is primarily a vehicle to deliver bug fixes. Any new features are minor changes resulting from fixes for performance, reliability, or usability issues.

Because 1.3.1 is the first 1.3.x release for CDH 4, if you are on CDH 4, also consult [New Features in Impala 1.3.0 / CDH 5.0.0](#) on page 1893 for more features that are new to you.



Note:

- The Impala 1.3.1 release is available for both CDH 4 and CDH 5. This is the first release in the 1.3.x series for CDH 4.

- A new `impalad` startup option, `--insert_inherit_permissions`, causes Impala `INSERT` statements to create each new partition with the same HDFS permissions as its parent directory. By default, `INSERT` statements create directories for new partitions using default HDFS permissions. See [INSERT Statement](#) for examples of `INSERT` statements for partitioned tables.
- The `SHOW FUNCTIONS` statement now displays the return type of each function, in addition to the types of its arguments. See [SHOW Statement](#) for examples.

- You can now specify the clause `FIELDS TERMINATED BY '\0'` with a `CREATE TABLE` statement to use text data files that use ASCII 0 (`\0`) characters as a delimiter. See [Using Text Data Files with Impala Tables](#) for details.
- In Impala 1.3.1 and higher, the `REGEXP` and `RLIKE` operators now match a regular expression string that occurs anywhere inside the target string, the same as if the regular expression was enclosed on each side by `.*`. See [REGEXP Operator](#) for examples. Previously, these operators only succeeded when the regular expression matched the entire target string. This change improves compatibility with the regular expression support for popular database systems. There is no change to the behavior of the `regexp_extract()` and `regexp_replace()` built-in functions.

New Features in Impala 1.3.0 / CDH 5.0.0



Note:

- The Impala 1.3.1 release is available for both CDH 4 and CDH 5. This is the first release in the 1.3.x series for CDH 4.

- The admission control feature lets you control and prioritize the volume and resource consumption of concurrent queries. This mechanism reduces spikes in resource usage, helping Impala to run alongside other kinds of workloads on a busy cluster. It also provides more user-friendly conflict resolution when multiple memory-intensive queries are submitted concurrently, avoiding resource contention that formerly resulted in out-of-memory errors. See [Admission Control and Query Queuing](#) for details.
- Enhanced `EXPLAIN` plans provide more detail in an easier-to-read format. Now there are four levels of verbosity: the `EXPLAIN_LEVEL` option can be set from 0 (most concise) to 3 (most verbose). See [EXPLAIN Statement](#) for syntax and [Understanding Impala Query Performance - EXPLAIN Plans and Query Profiles](#) for usage information.
- The `TIMESTAMP` data type accepts more kinds of input string formats through the `UNIX_TIMESTAMP` function, and produces more varieties of string formats through the `FROM_UNIXTIME` function. The documentation now also lists more functions for date arithmetic, used for adding and subtracting `INTERVAL` expressions from `TIMESTAMP` values. See [Impala Date and Time Functions](#) for details.
- New conditional functions, `NULLIF()`, `NULLIFZERO()`, and `ZEROIFNULL()`, simplify porting SQL containing vendor extensions to Impala. See [Impala Conditional Functions](#) for details.
- New utility function, `CURRENT_DATABASE()`. See [Impala Miscellaneous Functions](#) for details.
- Integration with the YARN resource management framework. Only available in combination with CDH 5. This feature makes use of the underlying YARN service, plus an additional service (Llama) that coordinates requests to YARN for Impala resources, so that the Impala query only proceeds when all requested resources are available. See [Resource Management for Impala](#) for full details.

On the Impala side, this feature involves some new startup options for the `impalad` daemon:

- `-enable_rm`
- `-llama_host`
- `-llama_port`
- `-llama_callback_port`
- `-cgroup_hierarchy_path`

For details of these startup options, see [Modifying Impala Startup Options](#).

This feature also involves several new or changed query options that you can set through the `impala-shell` interpreter and apply within a specific session:

- `MEM_LIMIT`: the function of this existing option changes when Impala resource management is enabled.
- `REQUEST_POOL`: a new option. (Renamed to `RESOURCE_POOL` in Impala 1.3.0.)
- `V_CPU_CORES`: a new option.
- `RESERVATION_REQUEST_TIMEOUT`: a new option.

For details of these query options, see [impala-shell Query Options for Resource Management](#).

New Features in Impala 1.2.4



Note: Impala 1.2.4 works with CDH 4. It is primarily a bug fix release for Impala 1.2.3, plus some performance enhancements for the catalog server to minimize startup and DDL wait times for Impala deployments with large numbers of databases, tables, and partitions.

- On Impala startup, the metadata loading and synchronization mechanism has been improved and optimized, to give more responsiveness when starting Impala on a system with a large number of databases, tables, or partitions. The initial metadata loading happens in the background, allowing queries to be run before the entire process is finished. When a query refers to a table whose metadata is not yet loaded, the query waits until the metadata for that table is loaded, and the load operation for that table is prioritized to happen first.
- Formerly, if you created a new table in Hive, you had to issue the `INVALIDATE METADATA` statement (with no table name) which was an expensive operation that reloaded metadata for all tables. Impala did not recognize the name of the Hive-created table, so you could not do `INVALIDATE METADATA new_table` to get the metadata for just that one table. Now, when you issue `INVALIDATE METADATA table_name`, Impala checks to see if that name represents a table created in Hive, and if so recognizes the new table and loads the metadata for it. Additionally, if the new table is in a database that was newly created in Hive, Impala also recognizes the new database.
- If you issue `INVALIDATE METADATA table_name` and the table has been dropped through Hive, Impala will recognize that the table no longer exists.
- New startup options let you control the parallelism of the metadata loading during startup for the `catalogd` daemon:
 - `--load_catalog_in_background` makes Impala load and cache metadata using background threads after startup. It is `true` by default. Previously, a system with a large number of databases, tables, or partitions could be unresponsive or even time out during startup.
 - `--num_metadata_loading_threads` determines how much parallelism Impala devotes to loading metadata in the background. The default is 16. You might increase this value for systems with huge numbers of databases, tables, or partitions. You might lower this value for busy systems that are CPU-constrained due to jobs from components other than Impala.

New Features in Impala 1.2.3



Note: Impala 1.2.3 works with CDH 4 and with CDH 5 beta 2. The resource management feature requires CDH 5 beta.

Impala 1.2.3 contains exactly the same feature set as Impala 1.2.2. Its only difference is one additional fix for compatibility with Parquet files generated outside of Impala by components such as Hive, Pig, or MapReduce.

New Features in Impala 1.2.2



Note: Impala 1.2.2 works with CDH 4. Its feature set is a superset of features in the Impala 1.2.0 beta, with the exception of resource management, which relies on CDH 5.

Impala 1.2.2 includes new features for performance, security, and flexibility. The major enhancements over 1.2.1 are performance related, primarily for join queries.

New user-visible features include:

- Join order optimizations. This highly valuable feature automatically distributes and parallelizes the work for a join query to minimize disk I/O and network traffic. The automatic optimization reduces the need to use query hints or to rewrite join queries with the tables in a specific order based on size or cardinality. The new `COMPUTE STATS` statement gathers statistical information about each table that is crucial for enabling the join optimizations. See [Performance Considerations for Join Queries](#) for details.
- `COMPUTE STATS` statement to collect both table statistics and column statistics with a single statement. Intended to be more comprehensive, efficient, and reliable than the corresponding Hive `ANALYZE TABLE` statement, which collects statistics in multiple phases through MapReduce jobs. These statistics are important for query planning for join queries, queries on partitioned tables, and other types of data-intensive operations. For optimal planning of join queries, you need to collect statistics for each table involved in the join. See [COMPUTE STATS Statement](#) for details.
- Reordering of tables in a join query can be overridden by the `STRAIGHT_JOIN` operator, allowing you to fine-tune the planning of the join query if necessary, by using the original technique of ordering the joined tables in descending order of size. See [Overriding Join Reordering with STRAIGHT_JOIN](#) for details.
- The `CROSS JOIN` clause in the `SELECT` statement to allow Cartesian products in queries, that is, joins without an equality comparison between columns in both tables. Because such queries must be carefully checked to avoid accidental overconsumption of memory, you must use the `CROSS JOIN` operator to explicitly select this kind of join. See [Cross Joins and Cartesian Products with the CROSS JOIN Operator](#) for examples.
- The `ALTER TABLE` statement has new clauses that let you fine-tune table statistics. You can use this technique as a less-expensive way to update specific statistics, in case the statistics become stale, or to experiment with the effects of different data distributions on query planning.
- LDAP username/password authentication in JDBC/ODBC. See [Enabling LDAP Authentication for Impala](#) for details.
- [GROUP_CONCAT\(\)](#) aggregate function to concatenate column values across all rows of a result set.
- The `INSERT` statement now accepts hints, `[SHUFFLE]` and `[NOSHUFFLE]`, to influence the way work is redistributed during `INSERT . . . SELECT` operations. The hints are primarily useful for inserting into partitioned Parquet tables, where using the `[SHUFFLE]` hint can avoid problems due to memory consumption and simultaneous open files in HDFS, by collecting all the new data for each partition on a specific node.
- Several built-in functions and operators are now overloaded for more numeric data types, to reduce the requirement to use `CAST()` for type coercion in `INSERT` statements. For example, the expression `2+2` in an `INSERT` statement formerly produced a `BIGINT` result, requiring a `CAST()` to be stored in an `INT` variable. Now, addition, subtraction, and multiplication only produce a result that is one step “bigger” than their arguments, and numeric and conditional functions can return `SMALLINT`, `FLOAT`, and other smaller types rather than always `BIGINT` or `DOUBLE`.
- New `fnv_hash()` built-in function for constructing hashed values. See [Impala Mathematical Functions](#) for details.
- The clause `STORED AS PARQUET` is accepted as an equivalent for `STORED AS PARQUETFILE`. This more concise form is recommended for new code.

Because Impala 1.2.2 builds on a number of features introduced in 1.2.1, if you are upgrading from an older 1.1.x release straight to 1.2.2, also review [New Features in Impala 1.2.1](#) on page 1895 to see features such as the `SHOW TABLE STATS` and `SHOW COLUMN STATS` statements, and user-defined functions (UDFs).

New Features in Impala 1.2.1



Note: Impala 1.2.1 works with CDH 4. Its feature set is a superset of features in the Impala 1.2.0 beta, with the exception of resource management, which relies on CDH 5.

Impala 1.2.1 includes new features for security, performance, and flexibility.

New user-visible features include:

- `SHOW TABLE STATS table_name` and `SHOW COLUMN STATS table_name` statements, to verify that statistics are available and to see the values used during query planning.
- `CREATE TABLE AS SELECT` syntax, to create a new table and transfer data into it in a single operation.
- `OFFSET` clause, for use with the `ORDER BY` and `LIMIT` clauses to produce “paged” result sets such as items 1-10, then 11-20, and so on.
- `NULLS FIRST` and `NULLS LAST` clauses to ensure consistent placement of `NULL` values in `ORDER BY` queries.
- New [built-in functions](#): `least()`, `greatest()`, `initcap()`.
- New aggregate function: `ndv()`, a fast alternative to `COUNT(DISTINCT col)` returning an approximate result.
- The `LIMIT` clause can now accept a numeric expression as an argument, rather than only a literal constant.
- The `SHOW CREATE TABLE` statement displays the end result of all the `CREATE TABLE` and `ALTER TABLE` statements for a particular table. You can use the output to produce a simplified setup script for a schema.
- The `--idle_query_timeout` and `--idle_session_timeout` options for `impalad` control the time intervals after which idle queries are cancelled, and idle sessions expire. See [Setting Timeout Periods for Daemons, Queries, and Sessions](#) for details.
- User-defined functions (UDFs). This feature lets you transform data in very flexible ways, which is important when using Impala as part of an ETL or ELT pipeline. Prior to Impala 1.2, using UDFs required switching into Hive. Impala 1.2 can run scalar UDFs and user-defined aggregate functions (UDAs). Impala can run high-performance functions written in C++, or you can reuse existing Hive functions written in Java.

You create UDFs through the `CREATE FUNCTION` statement and drop them through the `DROP FUNCTION` statement. See [User-Defined Functions \(UDFs\)](#) for instructions about coding, building, and deploying UDFs, and [CREATE FUNCTION Statement](#) and [DROP FUNCTION Statement](#) for related SQL syntax.

- A new service automatically propagates changes to table data and metadata made by one Impala node, sending the new or updated metadata to all the other Impala nodes. The automatic synchronization mechanism eliminates the need to use the `INVALIDATE METADATA` and `REFRESH` statements after issuing Impala statements such as `CREATE TABLE`, `ALTER TABLE`, `DROP TABLE`, `INSERT`, and `LOAD DATA`.

For even more precise synchronization, you can enable the [SYNC DDL](#) query option before issuing a DDL, `INSERT`, or `LOAD DATA` statement. This option causes the statement to wait, returning only after the catalog service has broadcast the applicable changes to all Impala nodes in the cluster.



Note:

Because the catalog service only monitors operations performed through Impala, `INVALIDATE METADATA` and `REFRESH` are still needed on the Impala side after creating new tables or loading data through the Hive shell or by manipulating data files directly in HDFS. Because the catalog service broadcasts the result of the `REFRESH` and `INVALIDATE METADATA` statements to all Impala nodes, when you do need to use those statements, you can do so a single time rather than on every Impala node.

This service is implemented by the `catalogd` daemon. See [The Impala Catalog Service](#) for details.

- `CREATE TABLE ... AS SELECT` syntax, to create a table and copy data into it in a single operation. See [CREATE TABLE Statement](#) for details.
- The `CREATE TABLE` and `ALTER TABLE` statements have new clauses `TBLPROPERTIES` and `WITH SERDEPROPERTIES`. The `TBLPROPERTIES` clause lets you associate arbitrary items of metadata with a particular table as key-value pairs. The `WITH SERDEPROPERTIES` clause lets you specify the serializer/deserializer (SerDes) classes that read and write data for a table; although Impala does not make use of these properties, sometimes particular values are needed for Hive compatibility. See [CREATE TABLE Statement](#) and [ALTER TABLE Statement](#) for details.

- Delegation support lets you authorize certain OS users associated with applications (for example, hue), to submit requests using the credentials of other users. Only available in combination with CDH 5. See [Configuring Impala Delegation for Hue and BI Tools](#) for details.
- Enhancements to `EXPLAIN` output. In particular, when you enable the new `EXPLAIN_LEVEL` query option, the `EXPLAIN` and `PROFILE` statements produce more verbose output showing estimated resource requirements and whether table and column statistics are available for the applicable tables and columns. See [EXPLAIN Statement](#) for details.
- `SHOW CREATE TABLE` summarizes the effects of the original `CREATE TABLE` statement and any subsequent `ALTER TABLE` statements, giving you a `CREATE TABLE` statement that will re-create the current structure and layout for a table.
- The `LIMIT` clause for queries now accepts an arithmetic expression, in addition to numeric literals.

New Features in Impala 1.2.0 (Beta)



Note: The Impala 1.2.0 beta release only works in combination with the beta version of CDH 5. The Impala 1.2.0 software is bundled together with the CDH 5 beta 1 download.

The Impala 1.2.0 beta includes new features for security, performance, and flexibility.

New user-visible features include:

- User-defined functions (UDFs). This feature lets you transform data in very flexible ways, which is important when using Impala as part of an ETL or ELT pipeline. Prior to Impala 1.2, using UDFs required switching into Hive. Impala 1.2 can run scalar UDFs and user-defined aggregate functions (UDAs). Impala can run high-performance functions written in C++, or you can reuse existing Hive functions written in Java.

You create UDFs through the `CREATE FUNCTION` statement and drop them through the `DROP FUNCTION` statement. See [User-Defined Functions \(UDFs\)](#) for instructions about coding, building, and deploying UDFs, and [CREATE FUNCTION Statement](#) and [DROP FUNCTION Statement](#) for related SQL syntax.

- A new service automatically propagates changes to table data and metadata made by one Impala node, sending the new or updated metadata to all the other Impala nodes. The automatic synchronization mechanism eliminates the need to use the `INVALIDATE METADATA` and `REFRESH` statements after issuing Impala statements such as `CREATE TABLE`, `ALTER TABLE`, `DROP TABLE`, `INSERT`, and `LOAD DATA`.



Note:

Because this service only monitors operations performed through Impala, `INVALIDATE METADATA` and `REFRESH` are still needed on the Impala side after creating new tables or loading data through the Hive shell or by manipulating data files directly in HDFS. Because the catalog service broadcasts the result of the `REFRESH` and `INVALIDATE METADATA` statements to all Impala nodes, when you do need to use those statements, you can do so a single time rather than on every Impala node.

This service is implemented by the `catalogd` daemon. See [The Impala Catalog Service](#) for details.

- Integration with the YARN resource management framework. Only available in combination with CDH 5. This feature makes use of the underlying YARN service, plus an additional service (Llama) that coordinates requests to YARN for Impala resources, so that the Impala query only proceeds when all requested resources are available. See [Resource Management for Impala](#) for full details.

On the Impala side, this feature involves some new startup options for the `impalad` daemon:

- `-enable_rm`
- `-llama_host`
- `-llama_port`

- `-llama_callback_port`
- `-cgroup_hierarchy_path`

For details of these startup options, see [Modifying Impala Startup Options](#).

This feature also involves several new or changed query options that you can set through the `impala-shell` interpreter and apply within a specific session:

- `MEM_LIMIT`: the function of this existing option changes when Impala resource management is enabled.
- `YARN_POOL`: a new option. (Renamed to `RESOURCE_POOL` in Impala 1.3.0.)
- `V_CPU_CORES`: a new option.
- `RESERVATION_REQUEST_TIMEOUT`: a new option.

For details of these query options, see [impala-shell Query Options for Resource Management](#).

- `CREATE TABLE ... AS SELECT` syntax, to create a table and copy data into it in a single operation. See [CREATE TABLE Statement](#) for details.
- The `CREATE TABLE` and `ALTER TABLE` statements have a new `TBLPROPERTIES` clause that lets you associate arbitrary items of metadata with a particular table as key-value pairs. See [CREATE TABLE Statement](#) and [ALTER TABLE Statement](#) for details.
- Delegation support lets you authorize certain OS users associated with applications (for example, `hue`), to submit requests using the credentials of other users. Only available in combination with CDH 5. See [Configuring Impala Delegation for Hue and BI Tools](#) for details.
- Enhancements to `EXPLAIN` output. In particular, when you enable the new `EXPLAIN_LEVEL` query option, the `EXPLAIN` and `PROFILE` statements produce more verbose output showing estimated resource requirements and whether table and column statistics are available for the applicable tables and columns. See [EXPLAIN Statement](#) for details.

New Features in Impala 1.1.1

Impala 1.1.1 includes new features for security and stability.

New user-visible features include:

- Additional security feature: auditing. New startup options for `impalad` let you capture information about Impala queries that succeed or are blocked due to insufficient privileges. To take full advantage of this feature with Cloudera Manager, upgrade to Cloudera Manager 4.7 or higher. For details, see [Impala Security Overview](#).
- Parquet data files generated by Impala 1.1.1 are now compatible with the Parquet support in Hive. See [Apache Impala Incompatible Changes and Limitations](#) on page 1925 for the procedure to update older Impala-created Parquet files to be compatible with the Hive Parquet support.
- Additional improvements to stability and resource utilization for Impala queries.
- Additional enhancements for compatibility with existing file formats.

New Features in Impala 1.1

Impala 1.1 includes new features for security, performance, and usability.

New user-visible features include:

- Extensive new security features, built on top of the Sentry open source project. Impala now supports fine-grained authorization based on roles. A policy file determines which privileges on which schema objects (servers, databases, tables, and HDFS paths) are available to users based on their membership in groups. By assigning privileges for views, you can control access to table data at the column level. For details, see [Impala Security Overview](#).
- Impala 1.1 works with Cloudera Manager 4.6 or higher. To use Cloudera Manager to manage authorization for the Impala web UI (the web pages served from port 25000 by default), use Cloudera Manager 4.6.2 or higher.
- Impala can now create, alter, drop, and query views. Views provide a flexible way to set up simple aliases for complex queries; hide query details from applications and users; and simplify maintenance as you rename or

reorganize databases, tables, and columns. See the overview section [Overview of Impala Views](#) and the statements [CREATE VIEW Statement](#), [ALTER VIEW Statement](#), and [DROP VIEW Statement](#).

- Performance is improved through a number of automatic optimizations. Resource consumption is also reduced for Impala queries. These improvements apply broadly across all kinds of workloads and file formats. The major areas of performance enhancement include:
 - Improved disk and thread scheduling, which applies to all queries.
 - Improved hash join and aggregation performance, which applies to queries with large build tables or a large number of groups.
 - Dictionary encoding with Parquet, which applies to Parquet tables with short string columns.
 - Improved performance on systems with SSDs, which applies to all queries and file formats.
- Some new built-in functions are implemented: [translate\(\)](#) to substitute characters within strings, [user\(\)](#) to check the login ID of the connected user.
- The new `WITH` clause for `SELECT` statements lets you simplify complicated queries in a way similar to creating a view. The effects of the `WITH` clause only last for the duration of one query, unlike views, which are persistent schema objects that can be used by multiple sessions or applications. See [WITH Clause](#).
- An enhancement to `DESCRIBE` statement, `DESCRIBE FORMATTED table_name`, displays more detailed information about the table. This information includes the file format, location, delimiter, ownership, external or internal, creation and access times, and partitions. The information is returned as a result set that can be interpreted and used by a management or monitoring application. See [DESCRIBE Statement](#).
- You can now insert a subset of columns for a table, with other columns being left as all `NULL` values. Or you can specify the columns in any order in the destination table, rather than having to match the order of the corresponding columns in the source. `VALUES` clause. This feature is known as “column permutation”. See [INSERT Statement](#).
- The new `LOAD DATA` statement lets you load data into a table directly from an HDFS data file. This technique lets you minimize the number of steps in your ETL process, and provides more flexibility. For example, you can bring data into an Impala table in one step. Formerly, you might have created an external table where the data files are not entirely under your control, or copied the data files to Impala data directories manually, or loaded the original data into one table and then used the `INSERT` statement to copy it to a new table with a different file format, partitioning scheme, and so on. See [LOAD DATA Statement](#).
- Improvements to Impala-HBase integration:
 - New query options for HBase performance: [HBASE_CACHE_BLOCKS](#) and [HBASE_CACHING](#).
 - Support for binary data types in HBase tables. See [Supported Data Types for HBase Columns](#) for details.
- You can issue `REFRESH` as a SQL statement through any of the programming interfaces that Impala supports. `REFRESH` formerly had to be issued as a command through the `impala-shell` interpreter, and was not available through a JDBC or ODBC API call. As part of this change, the functionality of the `REFRESH` statement is divided between two statements. In Impala 1.1, `REFRESH` requires a table name argument and immediately reloads the metadata; the new `INVALIDATE METADATA` statement works the same as the Impala 1.0 `REFRESH` did: the table name argument is optional, and the metadata for one or all tables is marked as stale, but not actually reloaded until the table is queried. When you create a new table in the Hive shell or through a different Impala node, you must enter `INVALIDATE METADATA` with no table parameter before you can see the new table in `impala-shell`. See [REFRESH Statement](#) and [INVALIDATE METADATA Statement](#).

New Features in Impala 1.0.1

The primary enhancements in Impala 1.0.1 are internal, for compatibility with the new Cloudera Manager 4.6 release. Try out the new **Impala Query Monitoring** feature in Cloudera Manager 4.6, which requires Impala 1.0.1.

New user-visible features include:

- The `VALUES` clause lets you `INSERT` one or more rows using literals, function return values, or other expressions. For performance and scalability, you should still use `INSERT ... SELECT` for bringing large quantities of data into an Impala table. The `VALUES` clause is a convenient way to set up small tables, particularly for initial testing of SQL features that do not require large amounts of data. See [VALUES Clause](#) for details.
- The `-B` and `-o` options of the `impala-shell` command can turn query results into delimited text files and store them in an output file. The plain text results are useful for using with other Hadoop components or Unix tools. In

benchmark tests, it is also faster to produce plain rather than pretty-printed results, and write to a file rather than to the screen, giving a more accurate picture of the actual query time.

- Several bug fixes. See [Issues Fixed in the 1.0.1 Release](#) on page 2457 for details.

New Features in Impala 1.0

This version has multiple performance improvements and adds the following functionality:

- Several bug fixes. See [Issues Fixed in the 1.0 GA Release](#) on page 2458.
- [ALTER TABLE](#) statement.
- [Hints](#) to allow specifying a particular join strategy.
- [REFRESH](#) for a single table.
- Dynamic resource management, allowing high concurrency for Impala queries.

New Features in Version 0.7 of the Impala Beta Release

This version has multiple performance improvements and adds the following functionality:

- Several bug fixes. See [Issues Fixed in Version 0.7 of the Beta Release](#) on page 2460.
- Support for the Parquet file format. For more information on file formats, see [How Impala Works with Hadoop File Formats](#).
- Added support for Avro.
- Support for the memory limits. For more information, see the example on modifying memory limits in [Modifying Impala Startup Options](#).
- Bigger and faster joins through the addition of partitioned joins to the already supported broadcast joins.
- Fully distributed aggregations.
- Fully distributed top-n computation.
- Support for creating and altering tables.
- Support for GROUP BY with floats and doubles.

In this version, both CDH 4.1 and 4.2 are supported, but due to performance improvements added, we highly recommend you use CDH 4.2 or higher to see the full benefit. If you are using Cloudera Manager, version 4.5 is required.

New Features in Version 0.6 of the Impala Beta Release

- Several bug fixes. See [Issues Fixed in Version 0.6 of the Beta Release](#) on page 2461.
- Added support for Impala on SUSE and Debian/Ubuntu. Impala is now supported on:
 - RHEL5.7/6.2 and Centos5.7/6.2
 - SUSE 11 with Service Pack 1 or higher
 - Ubuntu 10.04/12.04 and Debian 6.03
- Cloudera Manager 4.5 and CDH 4.2 support Impala 0.6.
- Support for the RCFile file format. For more information on file formats, see [Understanding File Formats](#).

New Features in Version 0.5 of the Impala Beta Release

- Several bug fixes. See [Issues Fixed in Version 0.5 of the Beta Release](#) on page 2462.
- Added support for a JDBC driver that allows you to access Impala from a Java client. To use this feature, follow the instructions in [Configuring Impala to Work with JDBC](#) to install the JDBC driver JARs on the client machine and modify the `CLASSPATH` on the client to include the JARs.

New Features in Version 0.4 of the Impala Beta Release

- Several bug fixes. See [Issues Fixed in Version 0.4 of the Beta Release](#) on page 2462.
- Added support for Impala on RHEL5.7/Centos5.7. Impala is now supported on RHEL5.7/6.2 and Centos5.7/6.2.
- Cloudera Manager 4.1.3 supports Impala 0.4.
- The Impala debug webserver now has the ability to serve static files from `${IMPALA_HOME}/www`. This can be disabled by setting `--enable_webserver_doc_root=false` on the command line. As a result, Impala now

uses the Twitter Bootstrap library to style its debug webpages, and the `/queries` page now tracks the last 25 queries run by each Impala daemon.

- Additional metrics available on the Impala Debug Webpage.

New Features in Version 0.3 of the Impala Beta Release

- Several bug fixes. See [Issues Fixed in Version 0.3 of the Beta Release](#) on page 2463.
- The `state-store-service` binary has been renamed `statestored`.
- The location of the Impala configuration files has changed from the `/usr/lib/impala/conf` directory to the `/etc/impala/conf` directory.

New Features in Version 0.2 of the Impala Beta Release

- Several bug fixes. See [Issues Fixed in Version 0.2 of the Beta Release](#) on page 2463.
- **Added Default Query Options** Default query options override all default QueryOption values when starting `impalad`. The format is:

```
-default_query_options='key=value;key=value'
```

What's New in Apache Kudu

This topic includes the release notes for all generally available versions (GA) of Apache Kudu.

Schema Design and Usage Limitations

The [Apache Kudu Known Issues](#) on page 2019 and [Usage Limitations](#) topics describes known issues and limitations with respect to schema design, integration with Apache Impala, and security in Kudu, as of the current release.

Kudu 1.7.0 / CDH 5.16.1 Release Notes

The release of Kudu in CDH 5.16.1 is based on Apache Kudu 1.7 with a few additional changes highlighted in these release notes.

Wire protocol compatibility, client library compatibility, incompatible changes and the issues stay the same as described in [Kudu 1.7.0 / CDH 5.15.0 Release Notes](#) on page 1902.

New Features in Kudu 1.7.0 / CDH 5.16.1

- The scan request timeout was added to Kudu scan token API.
- Two new configuration parameters have been added to the Spark connector to allow better tuning to avoid scan timeouts under high load: `socketReadTimeoutMs` and `scanRequestTimeout`
- The `kudu tablet unsafe_replace_tablet` tool was added to replace a tablet with a new one. Use this tool to recover a table when one of its tablets has permanently lost all replicas. The data in the replaced tablet is lost, so this tool should only be used as a last resort.
- Two new tools allow superusers to retrieve all the values of command line flags from remote Kudu processes: `kudu tserver get_flags` and `kudu master get_flags`. The `get_flag` tools support filtering the returned flags by tag. By default, they will return only flags that were explicitly set.
- A new CLI tool for manual data rebalancing can be used to redistribute table replicas among tablet servers. The rebalancing tool can be run via `kudu cluster rebalance` sub-command. Using the new tool, it's possible to rebalance Kudu clusters running CDH 5.12 (Kudu 1.4.0) and newer.

Optimizations and Improvements in Kudu 1.7.0 / CDH 5.16.1

- An optimization was added to improve the performance when scanning tables with large consecutive runs of deleted rows. For example, users may use `DELETE` all rows in a table or partition before re-adding them, or they may delete all data corresponding to some prefix of the primary key.
- The default timeout for starting tablet copy was increased from 3 to 30 seconds to make automatic tablet re-replication more robust when there is a heavy disk activity on the target server. Also, corresponding tablet server's flag `--tablet_copy_begin_session_timeout_ms` is now available as an advanced option.
- The `kudu cluster ksck` tool has been significantly enhanced.

- `ksck` now checks master health and consensus status, displays any unsafe or hidden flags set in the cluster, and produces a summary of the Kudu versions running on the master and tablet servers.
- `ksck` now supports JSON output, both in pretty-printed and compact form.
- The output format is controlled by the `-ksck_format` flag.
- The output of `kudu local_replica dump rowset` is now more human-readable.

Kudu 1.7.0 / CDH 5.15.0 Release Notes

New Features in Kudu 1.7.0 / CDH 5.15.0

- Kudu now supports the decimal column type. The decimal type is a numeric data type with fixed scale and precision suitable for financial and other arithmetic calculations where the imprecise representation and rounding behavior of float and double make those types impractical. The decimal type is also useful for integers larger than int64 and cases with fractional values in a primary key. See [Decimal Type](#) for more details.
- The strategy Kudu uses for automatically healing tablets which have lost a replica due to server or disk failures has been improved. The new re-replication strategy, or replica management scheme, first adds a replacement tablet replica before evicting the failed one. With the previous replica management scheme, the system first evicts the failed replica and then adds a replacement. The new replica management scheme allows for much faster recovery of tablets in scenarios where one tablet server goes down and then returns back shortly after 5 minutes or so. The new scheme also provides substantially better overall stability on clusters with frequent server failures. See [KUDU-1097](#) for more information.
- The `kudu fs update_dirs` tool now supports removing directories. Unless the `--force` flag is specified, Kudu will not allow the removal of a directory across which tablets are configured to spread data. If specified, all tablet replicas configured to use that directory will fail upon starting up and be replicated elsewhere, provided a majority exists elsewhere.
- Users can use the new `--fs_metadata_dir` to specify the directory in which to place tablet-specific metadata. It is recommended that this be placed on a high-performance drive with high bandwidth and low latency, such as a solid-state drive. If not specified, metadata will be placed in the directory specified by `--fs_wal_dir`, or in the directory specified by the first entry of `--fs_data_dirs` if metadata already exists there from an earlier deployment. Kudu does not automatically move existing metadata based on this configuration.
- `READ_YOUR_WRITES` is a new scan read mode. You can specify `READ_YOUR_WRITES` when creating a new scanner in C++, Java, and Python clients. If this mode is used, the client performs a read such that it follows all previously known writes and reads from this client. Reads in this mode ensure "read-your-writes" and "read-your-reads" session guarantees while minimizing latency caused by waiting for outstanding write transactions to complete. Note that this is still an experimental feature which may be stabilized in future releases.
- The tablet server web UI scans dashboard (`/scans`) has been improved with several new features:
 - Shows the most recently completed scans
 - Pseudo-SQL scan descriptor that concisely shows the selected columns and applied predicates
 - More complete, better documented scan statistics
- Kudu daemons now expose a web page `/stacks` which dumps the current stack trace of every thread running in the server. This information can be helpful when diagnosing performance issues.

Optimizations and Improvements in Kudu 1.7.0 / CDH 5.15.0

- By default, each tablet replica now stripes data blocks across three data directories instead of all data directories. This decreases the likelihood that any given tablet is affected in the event of a single disk failure. No substantial performance impact is expected due to this feature based on performance testing. This change only affects new replicas created after upgrading to Kudu 1.7.
- Kudu servers previously offered the ability to enable a separate metrics log which stores periodic snapshots of all metrics available on a server. This functionality is now available as part of a more general "diagnostics log" which is enabled by default. The diagnostics log includes periodic dumps of server metrics as well as collections of thread stack traces. The default configuration ensures that no more than 640MB of diagnostics logs are retained, and typically the space consumption is significantly less due to compression. The format and contents of this log file are documented in the [Apache Kudu Administration](#).

- The handling of errors in the synchronous Java client has been improved so that when an exception is thrown, the stack trace indicates the correct location where the client function was invoked rather than a call stack of an internal worker thread. The original call stack from the worker thread is available as a *suppressed exception*.
- The logging of errors in the Java client has been improved to exclude exception stack traces for expected scenarios such as failure to connect to a server in a cluster. Instead, only a single line informational message will be logged in such cases to aid in debugging.
- The Java client now uses a predefined prioritized list of TLS ciphers when establishing an encrypted connection to Kudu servers. This cipher list matches the list of ciphers preferred for server-to-server communication and ensures that the most efficient and secure ciphers are preferred. When the Kudu client is running on Java 8 or newer, this provides a substantial speed-up to read and write performance.
- Reporting for the `kudu cluster kscck` tool has been updated so tablets and tables with on-going tablet copies are shown as *recovering*. Additional reporting changes have been made to make various common scenarios, particularly tablet copies, less alarming.
- The performance of inserting rows containing many string or binary columns has been improved, especially in the case of highly concurrent write workloads.
- By default, Spark tasks that scan Kudu are now able to scan non-leader replicas. This allows Spark to more easily schedule `kudu-spark` tasks local to the data. You can disable this behavior by passing `leader_only` to the `kudu.scanLocality` option.
- The number of OS threads used in the steady state and during bursts of activity, such as in Raft leader elections triggered by a node failure, has been drastically reduced and should no longer exceed the value of `ulimit -u`. It should no longer be necessary to increase the value of `ulimit -u` or of `/proc/sys/kernel/threads-max` in order to run a Kudu tablet server in most cases. See [KUDU-1913](#) for more information.
- An issue where sparse column predicates could cause excessive data-block reads has been fixed. Previously in certain scans with sparsely matching predicates on multiple columns, Kudu read and decoded the same data blocks many times. This improvement typically results in a 5-10x performance increase for the affected scans. See [KUDU-2231](#) for more information.
- The efficiency and on-disk size of large updated values has been improved. This improves update-heavy workloads which overwrite large (1KiB+) values. See [KUDU-2253](#) for more information.

Wire Protocol Compatibility

Kudu 1.7.0 / CDH 5.15.0 is wire-compatible with previous versions of Kudu:

- Kudu 1.7 clients can connect to servers running Kudu 1.0 or later. If the client uses features that are not available on the target server, an error is returned.
- Rolling upgrade to Kudu 1.7 has not been sufficiently tested. We recommend against rolling upgrades in Kudu. You should shut down all nodes in the cluster, upgrade the software, and then restart the daemons on the new version.
- Kudu 1.0 clients may connect to servers running Kudu 1.7 with the exception of the below-mentioned restrictions regarding secure clusters.

The authentication features introduced in Kudu 1.3 place the following limitations on wire compatibility between Kudu 1.7 and versions earlier than 1.3:

- If a Kudu 1.7 cluster is configured with authentication or encryption is set to `required`, clients older than Kudu 1.3 are unable to connect.
- If a Kudu 1.7 cluster is configured with authentication and encryption is set to `optional` or `disabled`, older clients are still able to connect.

Client Library Compatibility

- The Kudu 1.7 Java client library is API- and ABI-compatible with Kudu 1.6. Applications written against Kudu 1.6 compile and run against the Kudu 1.7 client library and vice-versa.
- The Kudu 1.7 C++ client is API- and ABI-forward-compatible with Kudu 1.6. Applications written and compiled against the Kudu 1.6 client library run without modification against the Kudu 1.7 client library. Applications written and compiled against the Kudu 1.7 client library run without modification against the Kudu 1.6 client library.
- The Kudu 1.7 Python client is API-compatible with Kudu 1.6. Applications written against Kudu 1.6 continue to run against the Kudu 1.7 client and vice-versa.

- Kudu 1.7 clients that attempt to create a table with a decimal column on a target server running Kudu 1.6 or earlier receive an error response. Similarly Kudu clients running Kudu 1.6 or earlier result in an error when attempting to access any table containing a decimal column.

Kudu 1.7.0 / CDH 5.15.0 Incompatible Changes

- The newly introduced replica management scheme is not compatible with the old scheme, so it is not possible to run pre-1.7 Kudu masters with 1.7 Kudu tablet servers or vice versa. This is a server-side incompatibility only and it does not affect client compatibility. Kudu clients of prior versions are compatible with upgraded Kudu clusters:
 - Kudu masters of 1.7 version do not register Kudu tablet servers of 1.6 and prior versions.
 - Kudu tablet servers of 1.7 version does not work with Kudu masters of 1.6 and prior versions.
- The format of the previously-optional metrics log has changed to include a human-readable timestamp on each line. The path of the log file has also changed with the word “diagnostics” replacing the word “metrics” in the file name. The metrics log has been optimized to only include those metrics which have changed in between successive samples, and to not include entity attributes such as tablet partition information in the log. See [KUDU-2297](#) for more information.

Kudu 1.6.0 / CDH 5.14.2 Release Notes

New Features in Kudu 1.6.0 / CDH 5.14.2

- Cloudera Manager now supports installing Kudu as part of the Express Wizard installation.
- Cloudera Manager has improved monitoring of Kudu node health. Cloudera Manager will indicate whether the health of the service is concerning or bad according to the number of master or tablet server roles that are not running.

Kudu 1.6.0 / CDH 5.14.0 Release Notes

Apache Kudu 1.6.0 / CDH 5.14.0 is the first Kudu release to be supported with Cloudera Manager 5.14.x and CDH 5.14.x. Any future Kudu 1.6.x maintenance releases will be supported by corresponding CDH and Cloudera Manager 5.14.x maintenance releases.

New Features in Kudu 1.6.0 / CDH 5.14.0

- The disk failure tolerance feature is enabled by default on tablet servers, and the capability was extended to handle data directory failures at runtime. In the event of a runtime disk failure, any tablets with data on a failed disk will be shut down and restarted on another tablet server. By default, tablets are striped across all available disks. Note that the first configured data directory and the WAL directory cannot currently tolerate disk failures. This will be further improved in future Kudu releases.
- Kudu servers can now adopt new data directories via the `kudu fs update_dirs` tool. The new directory will be used by new tablet replicas only. Removing directories is not yet supported ([KUDU-2202](#)).
- Kudu servers have two new flags to control the webui TLS/HTTPS settings: `--webserver_tls_ciphers` and `--webserver_tls_min_protocol`. These flags allow the advertised TLS ciphers and TLS protocol versions to be configured. Additionally, the webserver now excludes insecure legacy ciphers by default ([KUDU-2190](#)).

Optimization and Improvements in Kudu 1.6.0 / CDH 5.14.0

- Kudu servers can now tolerate short interruptions in NTP clock synchronization. NTP synchronization is still required when any Kudu daemon starts up. If NTP synchronization is not available, diagnostic information is now logged to help pinpoint the issue ([KUDU-1578](#)).
- Tablet server startup time has been improved significantly on servers containing large numbers of blocks.
- The log block manager now performs disk data deletion in batches. This optimization can significantly reduce the time taken to delete data on a tablet.
- The usage of sensitive data redaction flag has been changed. By setting the `--redact=log` flag, redaction will be disabled in the web UI but retained for server logs. Alternatively, `--redact=none` can be used to disable redaction completely.

- The Spark DataSource integration now can take advantage of scan locality for better scan performance. The scan will take place at the closest replica instead of going to the leader.
- Various optimizations were made to reduce the 99th percentile latency of writes on tablet servers. This can also improve throughput on certain write workloads, particularly on larger clusters.
- You can configure Kudu to ignore the system-wide `auth_to_local` mappings configured in `/etc/krb5.conf` by setting the configuration flag `--use_system_auth_to_local=false` ([KUDU-2198](#)).
- The performance of the compaction scheduler has been improved. In previous versions, certain types of time series workloads could cause compaction scheduling to take tens of seconds. These workloads now schedule compactions an order of magnitude more efficiently.
- The compaction scheduler has been improved to avoid running a compaction when the benefit of that compaction is extremely small.
- Tablet servers now consider the health of all replicas of a tablet before deciding to evict one. This can improve stability of the Kudu cluster after experiencing multiple simultaneous daemon failures ([KUDU-2048](#)).
- Several performance improvements have been made to the Kudu master, particularly in concurrency of clients opening tables. This should improve performance in highly concurrent workloads.
- The on-disk size metric for a tablet now includes all data and metadata. Previously, it excluded WAL segments and consensus metadata ([KUDU-1755](#)).
- Added a verbose mode for the `kudu cluster ksck` command to enable output of detailed information on the cluster's metadata, even when no errors are detected.

Wire Protocol Compatibility

Kudu 1.6.0 is wire-compatible with previous versions of Kudu:

- Kudu 1.6 clients may connect to servers running Kudu 1.0 or later. If the client uses features that are not available on the target server, an error will be returned.
- Rolling upgrade between Kudu 1.5 and Kudu 1.6 servers may be possible though has not been sufficiently tested. Users are encouraged to shut down all nodes in the cluster, upgrade the software, and then restart the daemons on the new version.
- Kudu 1.0 clients may connect to servers running Kudu 1.6 with the exception of the below-mentioned restrictions regarding secure clusters.

The authentication features introduced in Kudu 1.3 place the following limitations on wire compatibility between Kudu 1.6 and versions earlier than 1.3:

- If a Kudu 1.6 cluster is configured with authentication or encryption set to "required", clients older than Kudu 1.3 will be unable to connect.
- If a Kudu 1.6 cluster is configured with authentication and encryption set to "optional" or "disabled", older clients will still be able to connect.

Client Library Compatibility

- The Kudu 1.6 Java client library is API- and ABI-compatible with Kudu 1.5. Applications written against Kudu 1.5 will compile and run against the Kudu 1.6 client library and vice-versa.
- The Kudu 1.6 C++ client is API- and ABI-forward-compatible with Kudu 1.5. Applications written and compiled against the Kudu 1.5 client library will run without modification against the Kudu 1.6 client library. Applications written and compiled against the Kudu 1.6 client library will run without modification against the Kudu 1.5 client library.
- The Kudu 1.6 Python client is API-compatible with Kudu 1.5. Applications written against Kudu 1.5 will continue to run against the Kudu 1.6 client and vice-versa.

Kudu 1.6.0 / CDH 5.13.2 Release Notes

New Features in Kudu 1.6.0 / CDH 5.13.2

- **New Command-line Debugging Tool** - The `kudu fs list` command line tool replaces exploratory usages of `kudu fs dump` and `kudu local_replica dump`. Output is formatted with default pretty-printing with options to output in CSV and JSON for scripts. Results can easily be filtered to a specific table, tablet, column, rowset, or block using flags. The tool can output many different fields: `table`, `table-id`, `tablet-id`, `partition`, `rowset-id`, `block-id`, `block-kind`, `column`, `column-id`, `cfile-data-type`, `cfile-encoding`, `cfile-compression`, `cfile-num-values`, `cfile-size`, `cfile-incompatible-features`, `cfile-compatible-features`, `cfile-min-key`, `cfile-max-key`, and `cfile-delta-stats`.

Kudu 1.5.0 / CDH 5.13.0 Release Notes

Apache Kudu 1.5.0 / CDH 5.13.0 is the first Kudu release to be supported with Cloudera Manager 5.13.x and CDH 5.13.x. Any future Kudu 1.5.x maintenance releases will be supported by corresponding CDH and Cloudera Manager 5.13.x maintenance releases.

New Features in Kudu 1.5.0 / CDH 5.13.0

- **Integrated into CDH** - Starting with Apache Kudu 1.5.0 / CDH 5.13.0, Kudu is fully integrated into CDH and ships as part of the CDH parcel and packages. On a cluster managed by Cloudera Manager, Kudu will be installed as part of CDH and does not need to be installed separately.
- **Kudu Web UI Improvements** - Kudu server web UIs have a new configuration dashboard (`/config`) which provides a high level summary of important security configuration values, such as whether RPC authentication is required, or web server HTTPS encryption is enabled. Other types of configuration will be added in future releases.
- **New Command-line Tools** - The `kudu` command line tool has two new features:
 - `kudu tablet change_config move_replica` - This tool moves a tablet replica from one tablet server to another, under the condition that the tablet is healthy. An operator can use this tool to rebalance tablet replicas between tablet servers.
 - `kudu local_replica data_size` - This tool summarizes the space usage of a tablet, breaking it down by type of file, column, and rowset.

Optimizations and Improvements in Kudu 1.5.0 / CDH 5.13.0

- The log block manager now performs disk synchronization in batches. This optimization can significantly reduce the time taken to copy tablet data from one server to another; in one case tablet copy time is reduced by 35%. It also improves the general performance of flushes and compactions.
- A new feature referred to as "tombstoned voting" is added to the Raft consensus subsystem to allow tablet replicas in the `TABLET_DATA_TOMBSTONED` state to vote in tablet leader elections. This feature increases Kudu's stability and availability by improving the likelihood that Kudu will be able to self-heal in more edge-case scenarios, such as when tablet copy operations fail. See [KUDU-871](#) for details.
- The tablet on-disk size metric has been made more accurate. Previously, the metric included only REDO deltas; it now counts all deltas. Additionally, the metric includes the size of bloomfiles, ad hoc indexes, and the tablet superblock. WAL segments and consensus metadata are still not counted. The latter is very small compared to the size of data, but the former may be significant depending on the workload (this will be resolved in a future release).
- The number of threads used by the Kudu tablet server has been further reduced. Previously, each follower tablet replica used a dedicated thread to detect leader tablet replica failures, and each leader replica used one dedicated thread per follower to send Raft heartbeats to that follower. The work performed by these dedicated threads has been reassigned to other threads. Other improvements were made to facilitate better thread sharing by tablets. For the purpose of capacity planning, expect the Kudu tablet server to create one thread for every five "cold" (i.e. those not servicing writes) tablets, and an additional three threads for every "hot" tablet. This will be further improved upon in future Kudu releases.

Wire Protocol Compatibility

Kudu 1.5.0 is wire-compatible with previous versions of Kudu:

- Kudu 1.5 clients may connect to servers running Kudu 1.0 or later. If the client uses features that are not available on the target server, an error will be returned.
- Rolling upgrade between Kudu 1.4 and Kudu 1.5 servers is believed to be possible though has not been sufficiently tested. Users are encouraged to shut down all nodes in the cluster, upgrade the software, and then restart the daemons on the new version.
- Kudu 1.0 clients may connect to servers running Kudu 1.5 with the exception of the below-mentioned restrictions regarding secure clusters.

The authentication features introduced in Kudu 1.3 place the following limitations on wire compatibility between Kudu 1.5 and versions earlier than 1.3:

- If a Kudu 1.5 cluster is configured with authentication or encryption set to "required", clients older than Kudu 1.3 will be unable to connect.
- If a Kudu 1.5 cluster is configured with authentication and encryption set to "optional" or "disabled", older clients will still be able to connect.

Client Library Compatibility

- The Kudu 1.5 Java client library is API- and ABI-compatible with Kudu 1.4. Applications written against Kudu 1.4 will compile and run against the Kudu 1.5 client library and vice-versa, unless one of the following newly added APIs is used:
- The Kudu 1.5 C++ client is API- and ABI-forward-compatible with Kudu 1.4. Applications written and compiled against the Kudu 1.4 client library will run without modification against the Kudu 1.5 client library. Applications written and compiled against the Kudu 1.5 client library will run without modification against the Kudu 1.4 client library.
- The Kudu 1.5 Python client is API-compatible with Kudu 1.4. Applications written against Kudu 1.4 will continue to run against the Kudu 1.5 client and vice-versa.

Known Issues and Limitations in Kudu 1.5.0 / CDH 5.13.0

For a complete list of known issues for Apache Kudu, see [Apache Kudu Known Issues](#) on page 2019.

Kudu 1.4.0 / CDH 5.12.1 Release Notes

Apache Kudu 1.4.0 / CDH 5.12.1 is a bug-fix release which includes the following fixes:

- [KUDU-2085](#) - Fixed a bug that caused crashes when seeking past the end of prefix-encoded blocks.
- [KUDU-2087](#) - Fixed an issue where Kudu would fail to start when Kerberos was enabled in FreeIPA-configured deployments.
- [KUDU-2053](#) - Fixed a race condition in the Java `RequestTracker`.
- [KUDU-2049](#) - Fixed an issue where scans on RLE-encoded integer columns would sometimes cause CHECK failures due to the CHECK condition being too strict.
- [KUDU-2052](#) - Kudu now uses `XFS_IOC_UNRESVSP64_ioctl` to punch holes on xfs filesystems. This fixes an issue with slow startup times on xfs when hole punching was done via `fallocate()`.
- [KUDU-1956](#) - Kudu will no longer crash if faced with a race condition when selecting rowsets for compaction.

Kudu 1.4.0 / CDH 5.12.0 Release Notes

Apache Kudu 1.4.0 / CDH 5.12.0 release is the first Kudu release to be supported with Cloudera Manager 5.12.x and CDH 5.12.x. Any future Kudu 1.4.x maintenance releases will be supported by corresponding CDH and Cloudera Manager 5.12.x maintenance releases.

New Features in Kudu 1.4.0 / CDH 5.12.0

- The C++ and Java client libraries now support the ability to alter the storage attributes (for example, encoding and compression) and default value of existing columns. Additionally, it is now possible to rename a column which is part of a table's primary key.
- The C++ client library now includes an experimental `KuduPartitioner` API which can be used to efficiently map rows to their associated partitions and hosts. This can be used to achieve better locality or distribution of writes in client applications.
- The Java client library now supports enabling fault tolerance on scanners. Fault tolerant scanners are able to transparently recover from concurrent server crashes at the cost of some performance overhead. See the [Java API documentation](#) for more details on usage.
- Kudu 1.4.0 includes the optional ability to compute, store, and verify checksums on all pieces of data stored on a server. Prior versions only performed checksums on certain portions of the stored data. This feature is not enabled by default since it makes a backward-incompatible change to the on-disk formats and thus prevents downgrades.
- **Kudu Command Line Tools**
 - The `kudu` command line tool now includes a new advanced administrative command, `kudu remote_replica unsafe_change_config`. This command can be used to force a tablet to perform an unsafe change of its Raft replication configuration. This can then be used to recover from scenarios such as a loss of a majority of replicas, at the risk of losing edits.
 - A new `kudu fs check` command can be used to perform various offline consistency checks on the local on-disk storage of a Kudu Tablet Server or Master. In addition to detecting various inconsistencies or corruptions, it can also detect and remove data blocks that are no longer referenced by any tablet but were not fully removed from disk due to a crash or a bug in prior versions of Kudu.
 - The `kudu` command line tool can now be used to list the addresses and identifiers of the servers in the cluster using either `kudu master list` or `kudu tserver list`.

For a complete list of Kudu's command line tools and usage information, see the [Command Line Reference documentation](#).

Optimizations and Improvements in Kudu 1.4.0 / CDH 5.12.0

- `kudu cluster ksck` now detects and reports new classes of inconsistencies and issues. In particular, it is better able to detect cases where a configuration change such as a replica eviction or addition is pending but is unable to be committed. It also properly detects and reports cases where a tablet has no elected leader.
- The default size for Write Ahead Log (WAL) segments has been reduced from 64MB to 8MB. Additionally, in the case that all replicas of a tablet are fully up to date and data has been flushed from memory, servers will now retain only a single WAL segment rather than two. These changes are expected to reduce the average consumption of disk space on the configured WAL disk by 16x, as well as improve the startup speed of tablet servers by reducing the number and size of WAL segments that need to be re-read.
- The default on-disk storage system used by Kudu servers (Log Block Manager) has been improved to compact its metadata and remove dead containers. This compaction and garbage collection occurs only at startup. Thus, the first startup after upgrade is expected to be longer than usual, and subsequent restarts should be shorter.
- The usability of the Kudu web interfaces has been improved, particularly for cases where a server hosts many tablets or a table has many partitions. Pages that list tablets now include a top-level summary of tablet status and show the complete list under a toggleable section.
- The maintenance manager features improved utilization of the configured maintenance threads. Previously, maintenance work would only be scheduled a maximum of 4 times per second, but now maintenance work will be scheduled immediately whenever any configured thread is available. This can improve the throughput of write-heavy workloads.
- The maintenance manager aggressively schedules flushes of in-memory data when memory consumption crosses 60 percent of the configured process-wide memory limit. The backpressure mechanism which begins to throttle client writes was also adjusted to not begin throttling until memory consumption reaches 80 percent of the configured limit. These two changes together result in improved write throughput, more consistent latency, and fewer timeouts due to memory exhaustion.
- Many performance improvements were made to improve write performance. Applications which send large batches of writes to Kudu should see substantially improved throughput in Kudu 1.4.0.

- Several improvements were made to reduce the memory consumption of Kudu Tablet Servers which hold large volumes of data. The specific amount of memory saved varies depending on workload, but the expectation is that approximately 350MB of excess peak memory usage has been eliminated per TB of data stored.
- The number of threads used by the Kudu Tablet Server has been reduced. Previously, each tablet used a dedicated thread to append to its WAL. Those threads now automatically stop running if there is no activity on a given tablet for a short period of time.

Fixed Issues in Kudu 1.4.0 / CDH 5.12.0

- [KUDU-2020](#) - Fixed an issue where re-replication after a failure would proceed significantly slower than expected. This bug caused many tablets to be unnecessarily copied multiple times before successfully being considered re-replicated, resulting in significantly more network and IO bandwidth usage than expected. Mean time to recovery on clusters with large amounts of data is improved by up to 10x by this fix.
- [KUDU-1982](#) - Fixed an issue where the Java client would call `NetworkInterface.getByInetAddress` very often, causing performance problems particularly on Windows where this function can be quite slow.
- [KUDU-1755](#) - Improved the accuracy of the `on_disk_size` replica metrics to include the size consumed by bloom filters, primary key indexes, superbloc metadata, and delta files. Note that because the size metric is now more accurate, the reported values are expected to increase after upgrading to Kudu 1.4.0. This does not indicate that replicas are using more space after the upgrade; rather, it is now accurately reporting the amount of space that has always been used.
- [KUDU-1192](#) - Kudu servers will now periodically flush their log messages to disk even if no `WARNING`-level messages have been logged. This makes it easier to tail the logs to see progress output during normal startup.

Wire Protocol Compatibility in Kudu 1.4.0 / CDH 5.12.0

- Kudu 1.4.0 clients can connect to servers running Kudu 1.0.0 or later. If the client uses features that are not available on the target server, an error will be returned.
- Kudu 1.0.0 clients can connect to servers running Kudu 1.4.0 with the exception of the restrictions on secure clusters mentioned further on in this list.
- Rolling upgrade between Kudu 1.3.0 and Kudu 1.4.0 servers is believed to be possible though it has not been sufficiently tested. Users are encouraged to shut down all nodes in the cluster, upgrade the software, and then restart the daemons on the new version.
- The authentication features introduced in Kudu 1.3.0 place the following limitations on wire compatibility between Kudu 1.4.0 and versions earlier than 1.3.0:
 - If a Kudu 1.4.0 cluster is configured with authentication or encryption set to `required`, clients older than Kudu 1.3.0 will be unable to connect.
 - If a Kudu 1.4.0 cluster is configured with authentication and encryption set to `optional` or `disabled`, older clients will still be able to connect.

Incompatible Changes in Kudu 1.4.0 / CDH 5.12.0

Kudu servers, by default, will now only allow unencrypted or unauthenticated connections from trusted subnets, which are private networks (127.0.0.0/8, 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16, 169.254.0.0/16) and local subnets of all local network interfaces. Unencrypted or unauthenticated connections from publicly routable IPs will be rejected, even if encryption and authentication are not configured.

The trusted subnets can be configured using the `--trusted_subnets` flag, which can be set to IP blocks represented in CIDR notation separated by comma. Set it to '0.0.0.0/0' to allow unauthenticated connections from all remote IP addresses. However, if network access is not otherwise restricted by a firewall, malicious users may be able to gain unauthorized access. This can be mitigated if authentication and encryption are configured to be required.

Client Library Compatibility

- **Java** - The Kudu 1.4.0 Java client library is API- and ABI-compatible with Kudu 1.3.0. Applications written against Kudu 1.3.0 will compile and run against the Kudu 1.4.0 client library and vice-versa, unless one of the following newly-added APIs is used:
 - `[Async]KuduScannerBuilder.setFaultTolerant(...)`

- **New methods** in `AlterTableOptions`: `removeDefault`, `changeDefault`, `changeDesiredBlockSize`, `changeEncoding`, `changeCompressionAlgorithm`
 - `KuduClient.updateLastPropagatedTimestamp`
 - `KuduClient.getLastPropagatedTimestamp`
 - **New getters** in `PartialRow`: `getBoolean`, `getByte`, `getShort`, `getInt`, `getLong`, `getFloat`, `getDouble`, `getString`, `getBinaryCopy`, `getBinary`, `isNull`, `isSet`.
- **C++** - The Kudu 1.4.0 C++ client is API- and ABI-forward-compatible with Kudu 1.3.0. Applications written and compiled against the Kudu 1.3.0 client library will run without modification against the Kudu 1.4.0 client library. Applications written and compiled against the Kudu 1.4.0 client library will run without modification against the Kudu 1.3.0 client library unless they use one of the following new APIs:
 - `KuduPartitionerBuilder`
 - `KuduPartitioner`
 - `KuduScanner::SetRowFormatFlags` (unstable API)
 - `KuduScanBatch::direct_data`, `KuduScanBatch::indirect_data` (unstable API)
 - **Python** - The Kudu 1.4.0 Python client is API-compatible with Kudu 1.3.0. Applications written against Kudu 1.3.0 will continue to run against the Kudu 1.4.0 client and vice-versa.

Known Issues and Limitations in Kudu 1.4.0 / CDH 5.12.0

For a complete list of schema design and usage limitations for Apache Kudu, see [Apache Kudu Known Issues](#) on page 2019.

Kudu 1.3.0 / CDH 5.11.2 Release Notes

Apache Kudu 1.3.0 / CDH 5.11.2 is a bug-fix release which includes the following fixes:

- [KUDU-2053](#) - Fixed a race condition in the Java `RequestTracker`.
- [KUDU-2049](#) - Fixed an issue where scans on RLE-encoded integer columns would sometimes cause CHECK failures due to the CHECK condition being too strict.
- [KUDU-1963](#) - Fixed an issue where the Java client misdiagnoses an error and logs a `NullPointerException` when a connection is closed by client while a negotiation is in progress.
- [KUDU-1853](#) - Fixed an issue where data blocks could be orphaned after a failed tablet copy.

Kudu 1.3.0 / CDH 5.11.1 Release Notes

Apache Kudu 1.3.0 / CDH 5.11.1 is a bug-fix release which includes the following fixes:

- [KUDU-1999](#) - Fixed an issue where the Kudu Spark connector would fail to kinit with the principal and keytab provided to a job.
- [KUDU-1993](#) - Fixed a validation issue with grouped gflags.
- [KUDU-1981](#) - Fixed an issue where Kudu server components would fail to start on machines with fully-qualified domain names longer than 64 characters when security was enabled. This was due to hard-coded restrictions in the OpenSSL library.
- [KUDU-1607](#) - Fixed a case in which a tablet replica on a tablet server could retain blocks of data which prevented it from being fully deleted.
- [KUDU-1933](#) - Fixed an issue in which a tablet server would crash and fail to restart after a single tablet received more than two billion write operations.

Kudu 1.3.0 / CDH 5.11.0 Release Notes



Note: Apache Kudu 1.3.0 / CDH 5.11.0 release is the first Kudu release to be supported with Cloudera Manager 5.11.x and CDH 5.11.x. Any future Kudu 1.3.x maintenance releases will be supported by corresponding CDH and Cloudera Manager 5.11.x maintenance releases.

New Features in Kudu 1.3.0 / CDH 5.11.0

- The Kudu CSD is now included in Cloudera Manager. Manual installation of the Kudu CSD is not required.
- Kudu 1.3 supports strong authentication using Kerberos. This feature allows users to authenticate themselves to Kudu using Kerberos tickets, and also provides mutual authentication of servers using Kerberos credentials stored in keytabs.
- Added support for encryption of data on the network using Transport Layer Security (TLS). Kudu will now use TLS to encrypt all network traffic between clients and servers as well as any internal traffic among servers, with the exception of traffic determined to be within a localhost network connection. Encryption is enabled by default whenever it can be determined that both the client and server support the feature.
- Added support for coarse-grained service-level authorization checks for access to the cluster. The operator may set up lists of permitted users who may act as administrators and as clients of the cluster. Combined with the strong authentication feature described above, this can enable a secure environment for some use cases. Note that fine-grained access control, for example table-level or column-level, is not yet supported.
- A background task was added to tablet servers that removes historical versions of data that have fallen behind the configured data retention time. This reduces disk space usage in all workloads, but particularly in those with a higher volume of updates or upserts. After upgrading, you might see disk usage go down a lot.
- Kudu now incorporates Google Breakpad, a library that writes crash reports in the case of a server crash. These reports can be found by default under the `minidumps` directory, or within a configured log directory, and can be useful during bug diagnosis.
- As of Kudu 1.3, multi-word flags can use a dash '-' separator in lieu of the underscore '_' separator. For example, the `--memory_limit_hard_bytes` flag can now be specified as `--memory-limit-hard-bytes`, or even `--memory_limit-hard_bytes`.

Optimizations and Improvements in Kudu 1.3.0 / CDH 5.11.0

- Kudu servers will now change the file permissions of data directories and contained data files based on a new configuration flag `--umask`. As a result, after upgrading, permissions on disk may be more restrictive than in previous versions. The new default configuration improves data security.
- Kudu web UIs will now redact strings that might include sensitive user data. For example, the monitoring page, which shows in-progress scans, no longer includes the scanner predicate values. The tracing and RPC diagnostics endpoints no longer include contents of RPCs, which may include table data.
- By default, Kudu now reserves 1% of each configured data volume as free space. If a volume is seen to have less than 1% of disk space free, Kudu stops writing to that volume to avoid completely filling up the disk.
- The default encoding for numeric columns (`int`, `float`, and `double`) has been changed to `BIT_SHUFFLE`. The default encoding for `binary` and `string` columns has been changed to `DICT_ENCODING`. Dictionary encoding automatically falls back to the old default (`PLAIN`) when cardinality is too high to be effectively encoded. These new defaults match the default behavior of other storage mechanisms, such as Apache Parquet, and are likely to perform better out of the box.

Existing tables with `AUTO_ENCODING`-configured columns will only be encoded during `DiskRowSet` compactions, or if new data is written to them. New tables and columns will have the new defaults.

- Kudu now uses `LZ4` compression when writing its Write Ahead Log (WAL). This improves write performance and stability for many use cases.
- Kudu now uses `LZ4` compression when writing delta files. This can improve both read and write performance as well as save substantial disk usage, especially for workloads involving a high number of updates or upserts containing compressible data.

Existing delta files will not get compressed until they are compacted. New files on both old and new tables alike will be compressed.

- The Kudu API now supports the ability to express `IS NULL` and `IS NOT NULL` predicates on scanners. The Spark DataSource integration will take advantage of these new predicates when possible.

- Both C++ and Java clients have been optimized to prune partitions more effectively when performing scans using the `IN (...)` predicate.
- The exception messages produced by the Java client are now truncated to a maximum length of 32KB.

Fixed Issues in Kudu 1.3.0 / CDH 5.11.0

- [KUDU-1968](#) - Fixed an issue in which the tablet server would delete an incorrect set of data blocks after an aborted attempt to copy a tablet from another server. This would produce data loss in unrelated tablets.
- [KUDU-1962](#) - Fixed a `NullPointerException` in the Java client in the case that the Kudu master is overloaded at the time the client requests location information. This could cause client applications to hang indefinitely regardless of configured timeouts.
- [KUDU-1893](#) - Fixed a critical bug in which wrong results would be returned when evaluating predicates applied to columns added using the `ALTER TABLE` operation.
- [KUDU-1905](#) - Fixed an issue where Kudu would crash after reinserts that resulted in an empty change list. This occurred in cases where the primary key was composed of all columns.
- [KUDU-1899](#) - Fixed a crash that occurred after inserting a row with an empty string as the single-column primary key.
- [KUDU-1904](#) - Fixed a potential crash when performing random reads against a column using RLE encoding and containing long runs of NULL values.
- [KUDU-1856](#) - Fixed an issue in which disk space could be leaked by Kudu servers storing data on partitions using the XFS file system. Any leaked disk space will now be automatically recovered upon upgrade.
- [KUDU-1888](#), [KUDU-1906](#) - Fixed multiple issues in the Java client where operation callbacks would never be triggered, causing the client to hang.

Wire Protocol Compatibility

Kudu 1.3.0 is wire-compatible with previous versions of Kudu:

- Kudu 1.3.0 clients can connect to servers running Kudu 1.0.0. If the client uses features that are not available on the target server, an error will be returned.
- Kudu 1.0.0 clients can connect to servers running Kudu 1.3.0 with the exception of the below-mentioned restrictions regarding secure clusters.
- Rolling upgrade between Kudu 1.2.0 and Kudu 1.3.0 servers is believed to be possible though has not been sufficiently tested. Users are encouraged to shut down all nodes in the cluster, upgrade the software, and then restart the daemons on the new version.

Security - The authentication features introduced in Kudu 1.3.0 place the following limitations on wire compatibility with older versions:

- If a Kudu 1.3.0 cluster is configured with authentication or encryption set to `required`, older clients will not be able to connect.
- If a Kudu 1.3.0 cluster is configured with authentication and encryption set to `optional` or `disabled`, older clients will still be able to connect.

Incompatible Changes in Kudu 1.3.0 / CDH 5.11.0

- Due to storage format changes in Kudu 1.3, downgrading from Kudu 1.3 to earlier versions is not supported. After upgrading to Kudu 1.3, attempting to restart with an earlier version will result in an error.
- In order to support running MapReduce and Spark jobs on secure clusters, these frameworks now connect to the cluster at job submission time to retrieve authentication credentials which can later be used by the tasks to be spawned. This means that the process submitting jobs to Kudu clusters must have direct access to that cluster.

- The embedded web servers in Kudu processes now specify the `X-Frame-Options: DENY` HTTP header which prevents embedding of Kudu web pages in HTML `iframe` elements.

Client Library Compatibility

- **Java** - The Kudu 1.3.0 Java client library is API- and ABI-compatible with Kudu 1.2.0. Applications written against Kudu 1.2.0 will compile and run against the Kudu 1.3.0 client library and vice-versa, unless one of the following newly-added APIs is used:
 - `[Async]KuduClient.exportAuthenticationCredentials(...)` (unstable API)
 - `[Async]KuduClient.importAuthenticationCredentials(...)` (unstable API)
 - `[Async]KuduClient.getMasterAddressesAsString()`
 - `KuduPredicate.newIsNotNullPredicate()`
 - `KuduPredicate.newIsNullPredicate()`
- **C++** - The Kudu 1.3.0 C++ client is API- and ABI-forward-compatible with Kudu 1.2.0. Applications written and compiled against the Kudu 1.2.0 client library will run without modification against the Kudu 1.3.0 client library. Applications written and compiled against the Kudu 1.3.0 client library will run without modification against the Kudu 1.2.0 client library unless they use one of the following new APIs:
 - `kudu::DisableOpenSSLInitialization()`
 - `KuduClientBuilder::import_authentication_credentials(...)`
 - `KuduClient::ExportAuthenticationCredentials(...)`
 - `KuduClient::NewIsNotNullPredicate(...)`
 - `KuduClient::NewIsNullPredicate(...)`
- **Python** - The Kudu 1.3.0 Python client is API-compatible with Kudu 1.2.0. Applications written against Kudu 1.2.0 will continue to run against the Kudu 1.3.0 client and vice-versa.

Known Issues and Limitations in Kudu 1.3.0 / CDH 5.11.0

For a complete list of schema design and usage limitations for Apache Kudu, see [Apache Kudu Known Issues](#) on page 2019.

Kudu 1.2.0 / CDH 5.10.2 Release Notes

Apache Kudu 1.2.x / CDH 5.10.2 includes the following fixed issues.

- [KUDU-1933](#) - Fixed an issue that truncated the 64-bit log index in the OpId to 32 bits, causing overflow of the log index.
- [KUDU-1607](#) - Fixed an issue where Kudu could not delete failed tablets using the `DROP TABLE` command.
- [KUDU-1905](#) - Allow reinserts on tables when all columns are part of the primary key.
- [KUDU-1893](#) - Made a fix to avoid incorrect NULL results and ensure evaluation of predicates for columns added after table creation.

Kudu 1.2.0 / CDH 5.10.1 Release Notes

Apache Kudu 1.2.x / CDH 5.10.1 includes the following fixed issues.

- [KUDU-1904](#) - Fixed a bug where RLE columns with only NULL values would crash on scan.
- [KUDU-1899](#) - Fixed an issue where tablet servers would crash after inserting an empty string primary key ("").
- [KUDU-1851](#) - Fixed an issue with the Python client which would crash whenever a `TableAlterer` is instantiated directly.
- [KUDU-1852](#) - `KuduTableAlterer` will no longer crash when given `nullptr` range bound arguments.
- [KUDU-1821](#) - Improved warnings when the catalog manager starts.

Kudu 1.2.0 / CDH 5.10.0 Release Notes



Note: Apache Kudu 1.2.0 / CDH 5.10.0 release is the first Kudu release to be supported with Cloudera Manager 5.10.0 and CDH 5.10.0. Any future Kudu 1.2.0 maintenance releases will likely be supported by corresponding CDH and Cloudera Manager 5.10.x maintenance releases.

Apache Kudu 1.2.x / CDH 5.10.0 includes the following new features, fixed issues, and changes.

New Features and Improvements in Kudu 1.2.0 / CDH 5.10.0

See also [Issues resolved for Kudu 1.2.0](#) and [Git changes between 1.1.x and 1.2.x](#).

New Features

- Kudu clients and servers now redact user data such as cell values from log messages, Java exception messages, and Status strings. User metadata such as table names, column names, and partition bounds are not redacted.
- Kudu's ability to provide consistency guarantees has been substantially improved:
 - Replicas now correctly track their "safe timestamp". This timestamp is the maximum timestamp at which reads are guaranteed to be repeatable.
 - A scan created using the `SCAN_AT_SNAPSHOT` mode will now either wait for the requested snapshot to be "safe" at the replica being scanned, or be re-routed to a replica where the requested snapshot is "safe". This ensures that all such scans are repeatable.
 - Kudu Tablet Servers now properly retain historical data when a row with a given primary key is inserted and deleted, followed by the insertion of a new row with the same key. Previous versions of Kudu would not retain history in such situations. This allows the server to return correct results for snapshot scans with a timestamp in the past, even in the presence of such "reinsertion" scenarios.
 - The Kudu clients now automatically retain the timestamp of their latest successful read or write operation. Scans using the `READ_AT_SNAPSHOT` mode without a client-provided timestamp automatically assign a timestamp higher than the timestamp of their most recent write. Writes also propagate the timestamp, ensuring that sequences of operations with causal dependencies between them are assigned increasing timestamps. Together, these changes allow clients to achieve read-your-writes consistency, and also ensure that snapshot scans performed by other clients return causally-consistent results.
- User data in log files is now redacted by default.
- Kudu servers now automatically limit the number of log files being stored. By default, 10 log files will be retained at each severity level.

Optimizations and Improvements

- The logging in the Java and `cpp` clients has been substantially quieted. Clients no longer log messages in normal operation unless there is some kind of error.
- The `cpp` client now includes a `KuduSession::SetErrorBufferSize` API which can limit the amount of memory used to buffer errors from asynchronous operations.
- The Java client now fetches tablet locations from the Kudu Master in batches of 1000, increased from batches of 10 in prior versions. This can substantially improve the performance of Spark and Impala queries running against Kudu tables with large numbers of tablets.
- Table metadata lock contention in the Kudu Master was substantially reduced. This improves the performance of tablet location lookups on large clusters with a high degree of concurrency.
- Lock contention in the Kudu Tablet Server during high-concurrency write workloads was also reduced. This can reduce CPU consumption and improve performance when a large number of concurrent clients are writing to a smaller number of servers.
- Lock contention when writing log messages has been substantially reduced. This source of contention could cause high tail latencies on requests, and when under high load could contribute to cluster instability such as election storms and request timeouts.
- The `BITSHUFFLE` column encoding has been optimized to use the `AVX2` instruction set present on processors including Intel(R) Haswell and later. Scans on `BITSHUFFLE`-encoded columns are now up to 30% faster.

- The `kudu` tool now accepts hyphens as an alternative to underscores when specifying actions. For example, `kudu local-replica copy-from-remote` may be used as an alternative to `kudu local_replica copy_from_remote`.

Issues Fixed in Kudu 1.2.0 / CDH 5.10.0

See [Issues resolved for Kudu 1.2.0](#) and [Git changes between 1.1.x and 1.2.x](#).

- [KUDU-1508](#) - Fixed a long-standing issue in which running Kudu on `ext4` file systems could cause file system corruption. While this issue has been known to still manifest in certain rare cases, the corruption is harmless and can be repaired as part of a regular `fsck`. Switching from `ext4` to `xfss` will also solve the problem.
- [KUDU-1399](#) - Implemented an LRU cache for open files, which prevents running out of file descriptors on long-lived Kudu clusters. By default, Kudu will limit its file descriptor usage to half of its configured `ulimit`.
- [Gerrit #5192](#) - Fixed an issue which caused data corruption and crashes in the case that a table had a non-composite (single-column) primary key, and that column was specified to use `DICT_ENCODING` or `BITSHUFFLE` encodings. If a table with an affected schema was written in previous versions of Kudu, the corruption will not be automatically repaired; users are encouraged to re-insert such tables after upgrading to Kudu 1.2 or later.
- [Gerrit #5541](#) - Fixed a bug in the Spark `KuduRDD` implementation which could cause rows in the result set to be silently skipped in some cases.
- [KUDU-1551](#) - Fixed an issue in which the tablet server would crash on restart in the case that it had previously crashed during the process of allocating a new WAL segment.
- [KUDU-1764](#) - Fixed an issue where Kudu servers would leak approximately 16-32MB of disk space for every 10GB of data written to disk. After upgrading to Kudu 1.2 or later, any disk space leaked in previous versions will be automatically recovered on startup.
- [KUDU-1750](#) - Fixed an issue where the API to drop a range partition would drop any partition with a matching lower `_or_` upper bound, rather than any partition with matching lower `_and_` upper bound.
- [KUDU-1766](#) - Fixed an issue in the Java client where equality predicates which compared an integer column to its maximum possible value (e.g. `Integer.MAX_VALUE`) would return incorrect results.
- [KUDU-1780](#) - Fixed the `kudu-client` Java artifact to properly shade classes in the `com.google.thirdparty` namespace. The lack of proper shading in prior releases could cause conflicts with certain versions of Google Guava.
- [Gerrit #5327](#) - Fixed shading issues in the `kudu-flume-sink` Java artifact. The sink now expects that Hadoop dependencies are provided by Flume, and properly shades the Kudu client's dependencies.
- Fixed a few issues using the Python client library from Python 3.

Incompatible Changes in Kudu 1.2.0 / CDH 5.10.0

Apache Kudu 1.2.0 introduces the following incompatible changes:

- The replication factor of tables is now limited to a maximum of 7. In addition, it is no longer allowed to create a table with an even replication factor.
- The `GROUP_VARINT` encoding is now deprecated. Kudu servers have never supported this encoding, and now the client-side constant has been deprecated to match the server's capabilities.
- **Client Library Compatibility**
 - The Kudu 1.2 Java client is API- and ABI-compatible with Kudu 1.1. Applications written against Kudu 1.1 will compile and run against the Kudu 1.2 client and vice-versa.
 - The Kudu 1.2 `cpp` client is API- and ABI-forward-compatible with Kudu 1.1. Applications written and compiled against the Kudu 1.1 client will run without modification against the Kudu 1.2 client. Applications written and compiled against the Kudu 1.2 client will run without modification against the Kudu 1.1 client unless they use one of the following new APIs:
 - `kudu::DisableSaslInitialization()`
 - `KuduSession::SetErrorBufferSize(...)`
 - The Kudu 1.2 Python client is API-compatible with Kudu 1.1. Applications written against Kudu 1.1 will continue to run against the Kudu 1.2 client and vice-versa.

Known Issues and Limitations in Kudu 1.2.0 / CDH 5.10.0

- [KUDU-1893](#) - Certain query results incorrectly include rows with null values for predicates.
Solution - Upgrade to the latest maintenance release.
- **Schema and Usage Limitations** - For a complete list of schema design and usage limitations for Apache Kudu, see [Apache Kudu Known Issues](#) on page 2019.

Incompatible Changes and Limitations



Important:

For changes in operating system support and other major requirements, see [CDH and Cloudera Manager Supported Operating Systems](#) on page 20.

Apache Avro Incompatible Changes and Limitations

See [Apache Avro](#) on page 1848 section of [What's New In CDH 5.2.x](#) on page 1848 for two changes that could possibly affect you when you upgrade to CDH 5.2.0.

Apache Crunch Incompatible Changes and Limitations

The following changes introduced in CDH 5.2 are not backward compatible:

- The `MemPipeline` now checks to ensure that any `DoFns` that are passed to it are serializable. This is designed to catch non-serializable `DoFns` during testing.
- Scala's `Iterable` has been replaced by `TraversableOnce` inside `Scrunch flatMap` functions in order to support functions that return `Iterators`.

CDH 5.4.0 introduces new HBase APIs, which will probably require some changes to Crunch code developed against HBase 0.96 APIs. For more information, see the section on [Apache Crunch](#) on page 1839 under "What's New in CDH 5.4.0".

Apache DataFu Incompatible Changes and Limitations

- Upgraded from version 0.4 to 1.1.0 (this upgrade is not backwards compatible).
- Removed `ApplyQuantiles`, `AliasBagFields`.
- Renamed package `datafu.pig.numbers` to `datafu.pig.random`.
- Renamed package `datafu.pig.bag.sets` to `datafu.pig.sets`.
- Renamed `TimeCount` to `SessionCount`, moved to `datafu.pig.sessions`.

Apache Flume Incompatible Changes and Limitations

The CDH Flume version requires to be matched to a minimum Kafka version. This means that if you upgrade to Kafka 2.2, you need to upgrade brokers as well as clients.

- In CDH 5.7 and higher, the Flume connector to Kafka only works with Kafka 2.0 and higher.
- In CDH 5.14 and higher, the Flume connector to Kafka only works with Kafka 2.2 and higher.

We recommend that you have the Kafka broker running the same or newer version as the client.

Apache Hadoop Incompatible Changes and Limitations

Hadoop Common

The following incompatible changes have been introduced in CDH 5:

- [HADOOP-12672](#) introduced a change to the behavior of RPC timeouts. Because of this change, a `NoSuchMethodError` exception may occur if the `hadoop-common.jar` is present on the classpath with a newer version of the `hadoop-hdfs.jar`s. This can happen when deploying JARs to the distributed cache or reordering classpaths with flags, such as `mapreduce.job.user.classpath.first` or `hbase.use.dynamic.jars`.

Cloudera Issue: CDH-69462

Affected versions: Upgrades from CDH 5.11.0 to 5.11.1 or 5.11.2; upgrades from CDH 5.10.0 or 5.10.1 to 5.10.2

Workaround: Ensure that all client jars have the same maintenance release dot version for all Hadoop jars. Cloudera uses the following pattern for versions: major.minor.maintenance. For example, the maintenance version for CDH 5.11.2 is 2.

HDFS

The following incompatible changes have been introduced in CDH 5:

- [HADOOP-14014](#), [HADOOP-14814](#): API incompatibility introduced in the `FsServerDefaults` class
 - **Affected Versions:** CDH 5.12.0, 5.12.1, 5.12.2, 5.13.0, and 5.13.1
 - **Workaround:** Upgrade to CDH 5.13.2 or 5.14
- HDFS FUSE does not support ACL.
 - **Affected Versions:** All CDH 5 versions
 - **Cloudera Bug:** CDH-60847
- CDH-53734, [HDFS-7933](#):

In CDH5.10.0 and higher, the output of the `fsck` command changed to also contain information about decommissioned replicas, and therefore if you try to parse the output of `fsck`, the parser could fail.
- CDH-54432, [HADOOP-11385](#):

In CDH 5.12.0 and higher, `JMXJSONServlet` replaced JSONP output with Cross-origin resource sharing (CORS) to prevent potential cross-site scripting attacks. External applications relying on the existing output may need to be updated because the new output is not compatible.
- CDH-26136, [HDFS-6962](#): ACL inheritance conflicts with `umaskmode`.
 - **Affected Versions:** CDH5.5.6, CDH5.7.5, CDH5.7.6, CDH5.8.4, CDH5.8.5, and CDH 5.9.1 and higher
- CDH-42931, [HDFS-6434](#) Default permission for creating file should be 644 for `WebHdfs/HttpFS`.
 - **Affected Versions:** CDH5.5.5, CDH5.5.6, CDH5.7.3, CDH5.7.4, CDH5.7.5, CDH5.7.6, CDH5.8.2, CDH5.8.3, CDH5.8.4, CDH5.8.5, and CDH 5.9.0 and higher
- [HDFS-9085](#) Show renewer information in `DelegationTokenIdentifier#toString`.
 - **Affected Versions:** CDH5.4.11 CDH5.5.5 CDH5.5.6 CDH5.7.2 CDH5.7.3 CDH5.7.4 CDH5.7.5 CDH5.7.6 CDH5.8.2 CDH5.8.3 CDH5.8.4 CDH5.8.5, and CDH 5.9.0 and higher
- The `getSnapshottableDirListing()` method returns `null` when there are no snapshottable directories. This is a change from CDH 5 Beta 2 where the method returns an empty array instead.
- Files named `.snapshot` or `.reserved` must not exist within HDFS.
- [HADOOP-10020](#): Disable symlinks temporarily.
 - **Affected Versions:** All CDH 5 versions
- [HDFS-2832](#) - The HDFS internal layout version has changed between CDH 5 Beta 1 and CDH 5 Beta 2, so a file system upgrade is required to move an existing Beta 1 cluster to Beta 2.
- [HDFS-4451](#): HDFS balancer command returns exit code 0 on success instead of 1.
 - **Affected Versions:** All CDH 5 versions

- [HDFS-4594](#): WebHDFS open sets Content-Length header to what is specified by length parameter rather than how much data is actually returned.
 - **Affected Versions:** All CDH 5 versions
 - **Impact:** In CDH 5, Content-Length header will contain the number of bytes actually returned, rather than the request length.
- [HDFS-4659](#): Support setting execution bit for regular files.
 - **Affected Versions:** All CDH 5 versions
 - **Impact:** In CDH 5, files copied out of `copyToLocal` may now have the executable bit set if it was set when they were created or copied into HDFS.
- [HDFS-4997](#) - `libhdfs` functions now return correct error codes in `errno` in case of an error, instead of always returning 255.
 - **Affected Versions:** All CDH 5 versions
- [HDFS-5138](#) - The `-finalize` NameNode startup option has been removed. To finalize an in-progress upgrade, you should instead use the `hdfs dfsadmin -finalizeUpgrade` command while your NameNode is running, or while both NameNodes are running in a High Availability setup.
 - **Affected Versions:** All CDH 5 versions
- [HDFS-7279](#) - In CDH 5.5.0 and higher, DataNode WebHDFS implementation uses Netty as an HTTP server instead of Jetty.
- [HADOOP-13508](#) - In CDH 5.11 and higher, the behavior of `org.apache.hadoop.fs.permission.FsPermission#FsPermission(String mode)` changed to fix a bug in parsing sticky bits. The new behavior may cause incompatible changes if an application depends on the original behavior.
- [HDFS-11056](#) - CDH 5.5.6 and 5.9.1 fixed a critical block corruption bug, but the fix introduced a new bug. Specifically, DataNodes may accumulate too many open file descriptors for deleted meta files over time. The bug is fixed in newer versions, including CDH5.9.2 and CDH5.10.0
Workaround: Restart DataNodes.
- [HADOOP-14104](#) - In CDH 5.12.0 and higher, the client always asks the NameNode for the KMS provider path. When HDFS high availability is used and the nameservices share a name, Cloudera Manager's Backup and Disaster Recovery replication fails.
Workaround: Use unique names for the nameservices.
 - **Affected versions::** CDH 5.12.x, CDH 5.13.x, 5.14.x

Change in High-Availability Support

In CDH 5, the only high-availability (HA) implementation is Quorum-based storage; shared storage using NFS is no longer supported.

MapReduce



Important: There is no separate tarball for MRv1. Instead, the MRv1 binaries, examples, and other contents are delivered in the Hadoop tarball itself. The scripts for running MRv1 are in the `bin-mapreduce1` directory in the tarball, and the MRv1 examples are in the `examples-mapreduce1` directory. You need to do some additional configuration; follow the directions below.

To use MRv1 from a tarball installation, complete the following steps:

1. Extract the files from the tarball.



Note: In the steps that follow, *install_dir* is the name of the directory into which you extracted the files.

2. Create a symbolic link as follows:

```
ln -s install_dir/bin-mapreduce1 install_dir/share/hadoop/mapreduce1/bin
```

3. Create a second symbolic link as follows:

```
ln -s install_dir/etc/hadoop-mapreduce1 install_dir/share/hadoop/mapreduce1/conf
```

4. Set the HADOOP_HOME and HADOOP_CONF_DIR environment variables in your execution environment as follows:

```
$ export HADOOP_HOME=install_dir/share/hadoop/mapreduce1
$ export HADOOP_CONF_DIR=$HADOOP_HOME/conf
```

5. Copy your existing `start-dfs.sh` and `stop-dfs.sh` scripts to *install_dir*/bin-mapreduce1

6. For convenience, add *install_dir*/bin to the PATH variable in your execution environment.

Apache MapReduce 2.0 (YARN) Incompatible Changes

The following incompatible changes occurred for Apache MapReduce 2.0 (YARN) between CDH 4.x and CDH 5 Beta 2:

- The CATALINA_BASE variable no longer determines whether a component is configured for YARN or MRv1. Use the `alternatives` command instead, and make sure CATALINA_BASE is not set.
- [YARN-1288](#) - YARN Fair Scheduler ACL change. Root queue defaults to everybody, and other queues default to nobody.
- YARN High Availability configurations have changed. Configuration keys have been renamed among other changes.
- The YARN_HOME property has been changed to HADOOP_YARN_HOME.
- Note the following changes to configuration properties in `yarn-site.xml`:
 - The value of `yarn.nodemanager.aux-services` should be changed from `mapreduce.shuffle` to `mapreduce_shuffle`.
 - `yarn.nodemanager.aux-services.mapreduce.shuffle.class` has been renamed to `yarn.nodemanager.aux-services.mapreduce_shuffle.class`
 - `yarn.resourcemanager.resourcemanager.connect.max.wait.secs` has been renamed to `yarn.resourcemanager.connect.max-wait.secs`
 - `yarn.resourcemanager.resourcemanager.connect.retry_interval.secs` has been renamed to `yarn.resourcemanager.connect.retry-interval.secs`
 - `yarn.resourcemanager.am.max-retries` is renamed to `yarn.resourcemanager.am.max-attempts`
 - The YARN_HOME environment variable used in the `yarn.application.classpath` has been renamed to HADOOP_YARN_HOME. Make sure you include `$HADOOP_YARN_HOME/*,$HADOOP_YARN_HOME/lib/*` in the classpath.

Changed default location of container executor binary

The default location of container-executor binary and .cfg files was changed to `/var/lib/yarn-ce`. It used to be `/opt/cloudera/parcels/<CDH_parcel_version>`.

Apache MapReduce 2.0 (YARN) Limitations

YARN rolling log aggregation not supported

Cloudera does not support YARN rolling log aggregation.

DockerContainerExecutor not supported in YARN

Cloudera does not support DockerContainerExecutor in YARN.

Node Manager configuration for YARN

A Spark on YARN job launches (with a ProcessBuilder shell executable) a Sqoop job with:

```
yarn jar /data/home/.../infaLib/sqoop-1.4.6-client.jar import -libjars \
  file:///data/home/.../lib/avro-mapred-1.7.5-hadoop2.jar --connect
jdbc:oracle:thin:@*****:1521 \
  --username ***** -m 1 --as-avrodatafile --columns CUSTOMER_ID,ORDER_ID,... --table
ORDERS \
  --target-dir hdfs://0.0.0.0:8020/user/.../d9b49b82_0f2e_41d3_bb31_ad6ff28aa966 \
  --password-file *****
```

This command appears to work fine from the command line, using the default configuration files from `/etc/hadoop/conf`. However, when you launch it from the Spark application (with a ProcessBuilder shell exec), configuration from the Cloudera parcel directory is picked up instead, and the job fails with the following error:

```
Log Length: 88
Error: Could not find or load main class org.apache.hadoop.mapreduce.v2.app.MRAppMaster
```

Workaround

To correct this issue:

1. Open the default `etc/hadoop/conf/mapred-site.xml`.
2. Copy the excerpt with the `mapreduce.application.classpath` property.
3. Go to the YARN NodeManager configuration page in Cloudera Manager.
4. Paste the excerpt in the **NodeManager Advanced Configuration Snippet** field. For example:

```
<property>
  <name>mapreduce.application.classpath</name>
  <value>$HADOOP_CLIENT_CONF_DIR,$HADOOP_CONF_DIR,$HADOOP_COMMON_HOME/,
    $HADOOP_COMMON_HOME/lib/, $HADOOP_HDFS_HOME/, $HADOOP_HDFS_HOME/lib/,
    $HADOOP_YARN_HOME/, $HADOOP_YARN_HOME/lib/
  </value>
</property>
```

Apache HBase Incompatible Changes and Limitations

Compatibility Notes for CDH 5

This section contains information that is relevant for all releases within the CDH 5 family. See the sections below for information which pertains to specific releases within CDH 5. If you are upgrading through more than one version (for instance, from CDH 5.0 to CDH 5.2), read the sections for each version, as most of the information listed applies to the given version and newer releases.

General Notes

- FilterBase no longer implements `Writable`. This means that you do not need to implement `readFields()` and `write()` methods when writing your own custom fields. Instead, put this logic into the `toByteArray` and `parseFrom` methods. See [this page](#) for an example.
- The default number of retained cell versions is reduced from 3 to 1. To increase the number of versions, you can specify the `VERSIONS` option at table creation or by altering existing tables. Starting with CDH 5.2, you can specify a global default number of versions, which will be applied to all newly created tables where the number of versions is not otherwise specified, by setting `hbase.column.max.version` to the desired number of versions in `hbase-site.xml`.
- In CDH 5 prior to 5.1.3, a Put submitted with a `KeyValue, KeyValue.Type.Delete` does not delete the cell. In CDH 5.1.3, this behavior is changed, so that a Put submitted with a `KeyValue, KeyValue.Type.Delete` does delete the cell. This fix is provided in [HBASE-11788](#).
- Cloudera does not provide support for user-provided custom coprocessors of any kind.

Operator API Changes

- [HBASE-6553](#) - Removed Avro Gateway. This feature was less robust and not used as much as the Thrift gateways. It has been removed upstream.
- HBase provides a metrics framework based on JMX beans. Between HBase 0.94 and 0.96, the metrics framework underwent many changes. Some beans were added and removed, some metrics were moved from one bean to another, and some metrics were renamed or removed. Click [here](#) to download the CSV spreadsheet which provides a mapping.

User API Changes

- The HBase User API (Get, Put, Result, Scanner etc; see [Apache HBase API documentation](#)) has evolved and attempts have been made to make sure the HBase Clients are source code compatible and thus should recompile without needing any source code modifications. This cannot be guaranteed however, since with the conversion to ProtoBufs, some relatively obscure APIs have been removed. Rudimentary efforts have also been made to preserve recompile compatibility with advanced APIs such as Filters and Coprocessors. These advanced APIs are still evolving and our guarantees for API compatibility are weaker here.
- As of 0.96, the User API has been marked and all attempts at compatibility in future versions will be made. A version of the javadoc that only contains the User API can be found [here](#).
- Other changes to CDH 5 HBase that require the upgrade include:
 - [HBASE-8015](#): The HBase Namespaces feature has changed HBase HDFS file layout.
 - [HBASE-4451](#): Renamed ZooKeeper nodes.
 - [HBASE-8352](#): HBase snapshots are now saved to the `/<hbase>/ .hbase-snapshot` dir instead of the `/ .snapshot` dir. This should be handled before upgrading HDFS.
 - [HBASE-7660](#): Removed support for HFile V1. All internal HBase files in the HFile v1 format must be converted to the HFile v2 format.
 - [HBASE-6170/HBASE-8909](#) - The `hbase.regionserver.lease.period` configuration parameter has been deprecated. Use `hbase.client.scanner.timeout.period` instead.

Compatibility Notes for CDH 5.15

The Admin interface has added methods for querying maintenance status within the cluster around split and merges. Downstream users who implement their own version of this interface will need to update their source code to account for the new methods. Downstream users who merely make use of Admin instances returned from the HBase API are not impacted.

Compatibility Notes for CDH 5.13

As a part of providing independent timeout tuning for reads and writes in HBASE-15866, the Table interface has added getter and setter methods. Downstream users who implement their own version of this interface will need to update their source code to include these methods. Downstream users who merely make use of Table instances returned from the HBase API are not impacted.

Compatibility Notes for CDH 5.12

There is a new `getTime` metric, - `*_region_*_metric_getTime`, that shows the time spent in gets, in milliseconds.

Compatibility Notes for CDH 5.9

- The default RPC scheduler has been changed from 'deadline' to 'fifo'. To reenable 'deadline', set `hbase.ipc.server.callqueue.type` to `deadline` in the `hbase-site.xml` file.
- Apache HBase no longer includes XSS defense or encoding for filters. Due to licensing issues, HBase no longer includes a prior XSS defense nor an encoding for filters. Additionally, several dependencies have been removed. Downstream users relying on transitive inclusion of the following will need to directly rely on the appropriate dependency themselves: jsr305 (from the FindBugs project), Apache Commons Fileupload, nekohtml, beanshell core, Apache xml graphics, OWASP antisamy, OWASP esapi, Xalan, Apache Xerces, and Xom.

Compatibility Notes for CDH 5.8

- HBase now ensures the jsr305 implementation from the findbugs project is not included in its binary artifacts or the compile / runtime dependencies of its user facing modules. Downstream users that rely on this jar will need to update their dependencies.
- HBase no longer includes Xerces implementation jars that were previously included via transitive dependencies. Downstream users relying on HBase for these artifacts will need to update their dependencies.
- This issue reverts fixes designed to prevent malicious content from rendering in HBase's UIs. Specifically, these changes shipped in 1.1.4+ and 1.2.0+. They were removed due to licensing issues discovered in the dependencies they introduced. Their implementation and those dependencies have been removed from HBase! Removal of these dependencies is against the strict definition of our version compatibility guidelines. However, inclusion of non-Apache approved licenses cannot be tolerated. Implementation of these fixes using an Apache-appropriate means is tracked in [HBASE-16328](#).

Compatibility Notes for CDH 5.7

- Cloudera recommends not using the new advanced configuration option `hbase.regionserver.hostname`, added in HBase 1.2 (CDH 5.7.0), which allows you to specify a separate external-facing hostname for a RegionServer.

Compatibility Notes for CDH 5.4

- The ports used by Apache HBase 1.0 changed from the 600XX range to the 160XX range. HBase in CDH reverted the change, and continues to use the 600XX port range, to maintain compatibility.
- If you used visibility labels prior to CDH 5.4 and assigned superuser privileges to HBase users by adding the `system` label to their set of labels, these users will no longer be superusers in CDH 5.4. To be sure that cached credentials are cleared, use the HBase Shell command `clear_auths <username>`, for each affected user. To grant users superuser privileges, add them to the **HBase Superusers** group in Cloudera Manager, or add them to the `hbase.superuser` property in `hbase-site.xml`, and restart the HMaster.
- `HTrace` is experimental in CDH 5.4.0. Artifacts and package names cannot be relied upon.
- Jersey was updated from 1.8 to 1.9. This has the following implications.
 - The Jersey version is now consistent with Apache HBase and other CDH components.
 - If your project relies upon `jersey-server`, you may need to make modifications.
- Curator in Hadoop was updated from 2.6.0 to 2.7.1. This has the following implications for HBase.
 - `PathUtils.validatePath(String)` changed return types, which will cause runtime errors for code compiled against the older version.
 - The `SharedCountReader` and `SharedValueReader` interfaces each added a method, which will cause compilation errors for code made to use the old version.
- `commons-codec` was upgraded from 1.7 to 1.9. This has the following implications for HBase.
 - The class `org.apache.commons.codec.net.QuotedPrintableCodec` has a constructor that throws additional exceptions. See the [API reference](#) for details.
- `commons-logging` was updated from version 1.1.1. to 1.2. This has the following implications for HBase.
 - `org.apache.commons.logging.LogSource.setLogImplementation(String)` no longer throws `ExceptionInInitializerError`, which may change behavior of code that expects it.
- CDH reverted API changes in HBase 1.0 which broke compatibility with HBase in CDH 5.0, 5.1, 5.2, and 5.3. If you have written applications using Apache HBase 1.0 APIs, you may need to modify these applications to run in CDH 5.4.

Differences between CDH 5.4 HBase 1.0 and Apache HBase 1.0:

- CDH 5.4.0 keeps `commons-math` at version 2.1 to maintain compatibility with earlier CDH releases, whereas Apache HBase 1.0 uses `commons-math 2.2`.
- CDH 5.4.0 keeps Netty at version 3 to maintain compatibility with earlier CDH releases, whereas Apache HBase 1.0 uses Netty 4.

Compatibility Notes for CDH 5.3

- The `Put` class no longer implements `Writable`. Instead, you can change the definition to `WritableComparable` if you have only Puts, or `WritableComparableDeletable` if you have a mix of Puts, Gets, and Deletes.

Compatibility Notes for CDH 5.2

- In HBase in CDH 5.1, the default value for `hbase.security.access.early_out` was set to `false`. In CDH 5.2, the default value has been changed to `true`. When set to `true`, if a user is not granted access to a column family qualifier, the `AccessController` immediately throws an `AccessDeniedException`. This change to the default behavior will affect users who enabled HFile version 3 and the `AccessController` coprocessor in CDH 5.1, and then upgrade to CDH 5.2. In this case, if you prefer `hbase.security.access.early_out` to be disabled, explicitly set it to `false` in `hbase-site.xml`.
- Starting with CDH 5.2, you can specify a global default number of versions, which will be applied to all newly created tables where the number of versions is not otherwise specified, by setting `hbase.column.max.version` to the desired number of versions in `hbase-site.xml`.
- HBase in CDH 5.2 differs from Apache HBase 0.98.6 in that CDH does not include [HBASE-11546](#), which provides ZooKeeper-less region assignment. CDH omits this feature because it is an incompatible change that prevents an upgraded cluster from being rolled back to a previous version.

Developer Interface Changes

- HBase 0.98.5 removed `ClientSmallScanner` from the public API. HBase in CDH 5.2 restores the constructor to maintain backward compatibility, but in future releases of HBase, this class will no longer be public. You should change your code to use the `Scan.setSmall(true)` method instead.

Compatibility Notes for CDH 5.1

General Notes

- [HBASE-8218](#) changes `AggregationClient` by replacing the `byte[] tablename` parameters with `HTable table`. This means that coprocessors compiled against CDH 5.0.x won't run or compile in CDH 5.1 and later.
- In CDH 5.1 and later, `delete*` methods of the `Delete` class of the HBase Client API use the timestamp from the constructor, the same behavior as the `Put` class. (In previous versions, the `delete*` methods ignored the constructor's timestamp, and used the value of `HConstants.LATEST_TIMESTAMP`. This behavior was different from the behavior of the `add()` methods of the `Put` class.) See [HBASE-10964](#).
- In CDH 5.1 and newer, HBase introduces a new snapshot format ([HBASE-7987](#)). A snapshot created in HBase 0.98 cannot be read by HBase 0.96. HBase 0.98 can read snapshots produced in previous versions of HBase, and no conversion is necessary.
- In CDH 5.1, the default value for `hbase.security.access.early_out` was changed from `true` to `false`. A setting of `true` means that if a user is not granted access to a column family qualifier, the `AccessController` immediately throws an `AccessDeniedException`. *This behavior change was reverted for CDH 5.2.*

Developer Interface Changes

- `HTablePool` is no longer supported in CDH 5.1 and later. The `HConnection` object is the replacement. You create the connection once and pass it around, as with the old table pool.

```
HConnection connection = HConnectionManager.createConnection(config);
HTableInterface table = connection.getTable(tableName);
table.put(put);
table.close();
connection.close();
```


You can set the `hbase.hconnection.threads.max` property in `hbase-site.xml` to control the pool size or you can pass an `ExecutorService` to `HConnectionManager.createConnection()`.

```
ExecutorService pool = ...;
HConnection connection = HConnectionManager.createConnection(conf, pool);
```

Compatibility Notes for CDH 5 Beta Releases

**Warning:**

CDH 5 Beta 1 and Beta 2 are not intended for production use, and have been superseded by official releases in the CDH 5 family.

The HBase client from CDH 5 Beta 1 is not wire compatible with CDH 5 Beta 2 because of changes introduced in [HBASE-9612](#). As a consequence, CDH 5 Beta 1 users will not be able to execute a rolling upgrade to CDH 5 Beta 2 (or later). This patch unifies the way the HBase clients make requests and simplifies the internals, but breaks wire compatibility. Developers may need to recompile applications built upon the CDH 5 Beta 1 API.

As of CDH 5 Beta 1 (HBase 0.95), the value of `hbase.regionserver.checksum.verify` defaults to `true`; in earlier releases the default is `false`.

Compatibility between CDH Beta and Apache HBase Releases

- Apache HBase 0.95.2 is not wire compatible with CDH 5 Beta 1 HBase 0.95.2.
- Apache HBase 0.96.x should be wire compatible with CDH 5 Beta 2 HBase 0.96.1.1.

Apache Hive Incompatible Changes and Limitations



Note: As of CDH 5, HCatalog is part of Apache Hive; incompatible changes in HCatalog are included below.

Metastore schema upgrade: CDH 5.2.0 includes Hive version 0.13.1. Upgrading from an earlier Hive version to Hive 0.13.1 or later requires a metastore schema upgrade.

**Warning:**

You must upgrade the metastore schema before starting the new version of Hive. Failure to do so may result in metastore corruption.

CDH 5 includes a new offline tool called `schematool`; Cloudera recommends you use this tool to upgrade your metastore schema.

Hive upgrade: Upgrading Hive from an earlier CDH 5.x release to CDH 5.2 or later, requires several manual steps. Follow the upgrade guide closely.

Incompatible changes between any earlier CDH version and CDH 5.4.x:

- CDH 5.2.0 and later clients cannot communicate with CDH 5.1.x and earlier servers. This means that you must upgrade the server before the clients.
- As of CDH 5.2.0, `DESCRIBE DATABASE` returns additional fields: `owner_name` and `owner_type`. The command will continue to behave as expected if you identify the field you're interested in by its (string) name, but could produce unexpected results if you use a numeric index to identify the field(s).
- CDH 5.2.0 implements [HIVE-6248](#), which includes some backward-incompatible changes to the HCatalog API.
- The CDH 5.2 Hive JDBC driver is not wire-compatible with the CDH 5.1 version of HiveServer2. Make sure you upgrade Hive clients and all other Hive hosts in tandem: the server first, and then the clients.
- HiveServer 1 is deprecated as of CDH 5.3, and will be removed in a future release of CDH. Users of HiveServer 1 should upgrade to HiveServer 2 as soon as possible.

- `org.apache.hcatalog` is deprecated as of CDH 5.3. All client-facing classes were moved from `org.apache.hcatalog` to `org.apache.hive.hcatalog` as of CDH 5.0 and the deprecated classes in `org.apache.hcatalog` will be removed altogether in a future release. If you are still using `org.apache.hcatalog`, you should move to `org.apache.hive.hcatalog` immediately.
- **Date partition columns:** as of Hive version 13, implemented in CDH 5.2, Hive validates the format of dates in partition columns, if they are stored as dates. A partition column with a date in invalid form can neither be used nor dropped once you upgrade to CDH 5.2 or higher. To avoid this problem, do one of the following:
 - Fix any invalid dates before you upgrade. Hive expects dates in partition columns to be in the form `YYYY-MM-DD`.
 - Store dates in partition columns as strings or integers.

You can use the following SQL query to find any partition-column values stored as dates:

```
SELECT "DBS"."NAME", "TBLS"."TBL_NAME", "PARTITION_KEY_VALS"."PART_KEY_VAL"
FROM "PARTITION_KEY_VALS"
INNER JOIN "PARTITIONS" ON "PARTITION_KEY_VALS"."PART_ID" = "PARTITIONS"."PART_ID"
INNER JOIN "PARTITION_KEYS" ON "PARTITION_KEYS"."TBL_ID" = "PARTITIONS"."TBL_ID"
INNER JOIN "TBLS" ON "TBLS"."TBL_ID" = "PARTITIONS"."TBL_ID"
INNER JOIN "DBS" ON "DBS"."DB_ID" = "TBLS"."DB_ID"
AND "PARTITION_KEYS"."INTEGER_IDX" = "PARTITION_KEY_VALS"."INTEGER_IDX"
AND "PARTITION_KEYS"."PKEY_TYPE" = 'date';
```

- **Decimal precision and scale:** As of CDH 5.4, Hive support for decimal precision and scale changes as follows:
 1. When `decimal` is used as a type, it means `decimal(10, 0)` rather than a precision of 38 with a variable scale.
 2. When Hive is unable to determine the precision and scale of a decimal type (for example in the case of non-generic User-Defined Function (UDF) that has an `evaluate()` method that returns `decimal`), a precision and scale of `(38, 18)` is assumed. In previous versions, a precision of 38 and a variable scale were assumed. Cloudera recommends you develop generic UDFs instead, and specify exact precision and scale.
 3. When a decimal value is assigned or cast to a different decimal type, rounding is used to handle cases in which the precision of the value is greater than that of the target decimal type, as long as the integer portion of the value can be preserved. In previous versions, if the value's precision was greater than 38 (the only allowed precision for the `decimal` type), the value was set to null, regardless of whether the integer portion could be preserved.
- **Deprecation of HivePassThrough serde formats:** As of CDH 5.4, [HIVE-8910](#) changes how the storage handler uses the `HivePassThroughOutputFormat` class. It removes the empty default constructor, which breaks `org.apache.hadoop.util.ReflectionUtils.newInstance` and throws a `NoSuchMethodException`. The workaround is to re-create the Hive tables without HivePassThrough serde formats.

Hue Incompatible Changes and Limitations

- You will need to upgrade any custom applications after you upgrade to CDH 5.4.0.
- All CDH 5 versions of Hue work only with the default system Python version of the operating system it is being installed on. For example, on RHEL/CentOS 6, you need Python 2.6 to start Hue.



Note: RHEL 5 and CentOS 5 users will have to download Python 2.6 from the EPEL repository.

Apache Impala Incompatible Changes and Limitations

The Impala version covered by this documentation library contains the following incompatible changes. These are things such as file format changes, removed features, or changes to implementation, default configuration, dependencies, or prerequisites that could cause issues during or after an Impala upgrade.

Even added SQL statements or clauses can produce incompatibilities, if you have databases, tables, or columns whose names conflict with the new keywords.

Incompatible Changes Introduced in CDH 5.14.x / Impala 2.11.x

The following new or changed behaviors could cause issues during upgrades or cause changed behavior from prior releases, possibly requiring changes to source code:

- In the [impala] section of the `.impalarc` file, you can specify 0 or 1 wherever a Boolean value is required. Under the [impala.query_options] section, values that are expected to be Boolean must be `true`, `false`, 0, or 1, otherwise Impala reports an error. Formerly, unrecognized Boolean values would only cause a warning.
- Arithmetic expressions involving both `DECIMAL` and `FLOAT` or `DOUBLE` arguments now produce results of `DECIMAL` type instead of `DOUBLE`.

For the full list of issues closed in this release, including any that introduce behavior changes or incompatibilities, see the [changelog for CDH 5.14](#).

Incompatible Changes Introduced in CDH 5.13.x / Impala 2.10.x

The following new or changed behaviors could cause issues during upgrades or cause changed behavior from prior releases, possibly requiring changes to source code:

- The buffer pool feature changes the way memory is allocated during a query, and the thresholds at which operators within a query might spill to disk. Although in general the result is less memory usage and less spilling, it is possible that some queries that worked before (even if those queries were slow and inefficient) might now fail, requiring you to adjust query options related to the buffer pool. In particular, if a table contains rows whose total data size is greater than 512 KB, you must increase the setting for the `MAX_ROW_SIZE` query option.

After upgrading to CDH 5.13 / Impala 2.10, follow the instructions in [Handling Large Rows During Upgrade to CDH 5.13 / Impala 2.10 or Higher](#) to check if your queries are affected by these changes and to modify your configuration settings if so. This advice is especially important for any users who have increased the `--read_size` configuration setting from its default value of 8 MB.

For the full list of issues closed in this release, including any that introduce behavior changes or incompatibilities, see the [changelog for CDH 5.13](#).

Incompatible Changes Introduced in CDH 5.12.x / Impala 2.9.x

For the full list of issues closed in this release, including any that introduce behavior changes or incompatibilities, see the [changelog for CDH 5.12](#).

Impala Incompatible Changes Introduced in CDH 5.11.x / Impala 2.9.x

Because the core Impala functionality in CDH 5.11.x is very close to that of CDH 5.10.x, there are no new incompatible changes for this release.

Impala Incompatible Changes Introduced in CDH 5.10.x / Impala 2.8.x

- Llama support is removed completely from Impala. Related flags (`--enable_rm`) and query options (such as `V_CPU_CORES`) remain but do not have any effect.
If `--enable_rm` is passed to Impala, a warning is printed to the log on startup.
- The syntax related to Kudu tables includes a number of new reserved words, such as `COMPRESSION`, `DEFAULT`, and `ENCODING`, that might conflict with names of existing tables, columns, or other identifiers from older Impala versions. See [Impala Reserved Words](#) for the full list of reserved words.
- The DDL syntax for Kudu tables, particularly in the `CREATE TABLE` statement, is different from the special `impala_next` fork that was previously used for accessing Kudu tables from Impala:
 - The `DISTRIBUTE BY` clause is now `PARTITIONED BY`.
 - The `INTO N BUCKETS` clause is now `PARTITIONS N`.
 - The `SPLIT ROWS` clause is replaced by different syntax for specifying the ranges covered by each partition.

- The `DESCRIBE` output for Kudu tables includes several extra columns.
- Non-primary-key columns can contain `NULL` values by default. The `SHOW CREATE TABLE` output for these columns displays the `NULL` attribute. There was a period during early experimental versions of Impala + Kudu where non-primary-key columns had the `NOT NULL` attribute by default.
- The `IGNORE` keyword that was present in early experimental versions of Impala + Kudu is no longer present. The behavior of the `IGNORE` keyword is now the default: DML statements continue with warnings, instead of failing with errors, if they encounter conditions such as “primary key already exists” for an `INSERT` statement or “primary key already deleted” for a `DELETE` statement.
- The replication factor for Kudu tables must be an odd number.
- A UDF compiled into an LLVM IR bitcode module (`.bc`) might encounter a runtime error when native code generation is turned off by setting the query option `DISABLE_CODEGEN=1`. This issue also applies when running a built-in or native UDF with more than 20 arguments. See [IMPALA-4432](#) for details. As a workaround, either turn native code generation back on with the query option `DISABLE_CODEGEN=0`, or use the regular UDF compilation path that does not produce an IR module.

Impala Incompatible Changes Introduced in CDH 5.9.x / Impala 2.7.x

- Bug fixes related to parsing of floating-point values (IMPALA-1731 and IMPALA-3868) can change the results of casting strings that represent invalid floating-point values. For example, formerly a string value beginning or ending with `inf`, such as `1.23inf` or `infinite`, now are converted to `NULL` when interpreted as a floating-point value. Formerly, they were interpreted as the special “infinity” value when converting from string to floating-point. Similarly, now only the string `NaN` (case-sensitive) is interpreted as the special “not a number” value. String values containing multiple dots, such as `3. .141` or `3.1.4.1`, are now interpreted as `NULL` rather than being converted to valid floating-point values.

Impala Incompatible Changes Introduced in CDH 5.8.x / Impala 2.6.x

- The default for the `RUNTIME_FILTER_MODE` query option is changed to `GLOBAL` (the highest setting).
- The `RUNTIME_BLOOM_FILTER_SIZE` setting is now only used as a fallback if statistics are not available; otherwise, Impala uses the statistics to estimate the appropriate size to use for each filter.
- Admission control and dynamic resource pools are enabled by default. When upgrading from an earlier release, you must turn on these settings yourself if they are not already enabled. See [Admission Control and Query Queuing](#) for details about admission control.
- Impala reserves some new keywords, in preparation for support for Kudu syntax: `buckets`, `delete`, `distribute`, `hash`, `ignore`, `split`, and `update`.
- For Kerberized clusters, the Catalog service now uses the Kerberos principal instead of the operating system user that runs the `catalogd` daemon. This eliminates the requirement to configure a `hadoop.user.group.static.mapping.overrides` setting to put the OS user into the Sentry administrative group, on clusters where the principal and the OS user name for this user are different.
- The mechanism for interpreting `DECIMAL` literals is improved, no longer going through an intermediate conversion step to `DOUBLE`:
 - Casting a `DECIMAL` value to `TIMESTAMP DOUBLE` produces a more precise value for the `TIMESTAMP` than formerly.
 - Certain function calls involving `DECIMAL` literals now succeed, when formerly they failed due to lack of a function signature with a `DOUBLE` argument.
- Improved type accuracy for `CASE` return values. If all `WHEN` clauses of the `CASE` expression are of `CHAR` type, the final result is also `CHAR` instead of being converted to `STRING`.

- The initial release of CDH 5.7 / Impala 2.5 sometimes has a higher peak memory usage than in previous releases while reading Parquet files. The following query options might help to reduce memory consumption in the Parquet scanner:
 - Reduce the number of scanner threads, for example: `set num_scanner_threads=30`
 - Reduce the batch size, for example: `set batch_size=512`
 - Increase the memory limit, for example: `set mem_limit=64g`

You can track the status of the fix for this issue at [IMPALA-3662](#).

- The `S3_SKIP_INSERT_STAGING` query option, which is enabled by default, increases the speed of `INSERT` operations for S3 tables. The speedup applies to regular `INSERT`, but not `INSERT OVERWRITE`. The tradeoff is the possibility of inconsistent output files left behind if a node fails during `INSERT` execution. See [S3_SKIP_INSERT_STAGING Query Option \(CDH 5.8 or higher only\)](#) for details.

Certain features are turned off by default, to avoid regressions or unexpected behavior following an upgrade. Consider turning on these features after suitable testing:

- Impala now recognizes the `auth_to_local` setting, specified through the HDFS configuration setting `hadoop.security.auth_to_local`. This feature is disabled by default; to enable it, specify `--load_auth_to_local_rules=true` in the `impalad` configuration settings.
- A new query option, `PARQUET_ANNOTATE_STRINGS_UTF8`, makes Impala include the UTF-8 annotation metadata for `STRING`, `CHAR`, and `VARCHAR` columns in Parquet files created by `INSERT` or `CREATE TABLE AS SELECT` statements.
- A new query option, `PARQUET_FALLBACK_SCHEMA_RESOLUTION`, lets Impala locate columns within Parquet files based on column name rather than ordinal position. This enhancement improves interoperability with applications that write Parquet files with a different order or subset of columns than are used in the Impala table.

Impala Incompatible Changes Introduced in CDH 5.7.x / Impala 2.5.x

- The admission control default limit for concurrent queries (the `max requests` setting) is now unlimited instead of 200.
- Multiplying a mixture of `DECIMAL` and `FLOAT` or `DOUBLE` values now returns `DOUBLE` rather than `DECIMAL`. This change avoids some cases where an intermediate value would underflow or overflow and become `NULL` unexpectedly. The results of multiplying `DECIMAL` and `FLOAT` or `DOUBLE` might now be slightly less precise than before. Previously, the intermediate types and thus the final result depended on the exact order of the values of different types being multiplied, which made the final result values difficult to reason about.
- Previously, the `_` and `%` wildcard characters for the `LIKE` operator would not match characters on the second or subsequent lines of multi-line string values. The fix for issue [IMPALA-2204](#) causes the wildcard matching to apply to the entire string for values containing embedded `\n` characters. This could cause different results than in previous Impala releases for identical queries on identical data.
- Formerly, all Impala UDFs and UDAs required running the `CREATE FUNCTION` statements to re-create them after each `catalogd` restart. In CDH 5.7 / Impala 2.5 and higher, functions written in C++ are persisted across restarts, and the requirement to re-create functions only applies to functions written in Java. Adapt any function-reloading logic that you have added to your Impala environment.
- `CREATE TABLE LIKE` no longer inherits HDFS caching settings from the source table.
- The `SHOW DATABASES` statement now returns two columns rather than one. The second column includes the associated comment string, if any, for each database. Adjust any application code that examines the list of databases and assumes the result set contains only a single column.

- The output of the `SHOW FUNCTIONS` statement includes two new columns, showing the kind of the function (for example, `BUILTIN`) and whether or not the function persists across catalog server restarts. For example, the `SHOW FUNCTIONS` output for the `_impala_builtins` database starts with:

```

+-----+-----+-----+-----+
| return type | signature | binary type | is |
| persistent | | | |
+-----+-----+-----+-----+
| BIGINT | abs(BIGINT) | BUILTIN | true |
| DECIMAL(*,*) | abs(DECIMAL(*,*)) | BUILTIN | true |
| DOUBLE | abs(DOUBLE) | BUILTIN | true |
| ... | | | |

```

Impala Incompatible Changes Introduced in CDH 5.6.x / Impala 2.4.x

The Impala feature set for CDH 5.6 is the same as for CDH 5.5. Therefore, there are no incompatible changes for Impala introduced in CDH 5.6.

Impala Incompatible Changes Introduced in CDH 5.5.x / Impala 2.3.x



Note:

The use of the Llama component for integrated resource management within YARN is no longer supported with CDH 5.5 / Impala 2.3 and higher. The Llama support code is removed entirely in CDH 5.10 / Impala 2.8 and higher.

For clusters running Impala alongside other data management components, you define static service pools to define the resources available to Impala and other components. Then within the area allocated for Impala, you can create dynamic service pools, each with its own settings for the Impala admission control feature.

- If Impala encounters a Parquet file that is invalid because of an incorrect magic number, the query skips the file. This change is caused by the fix for issue [IMPALA-2130](#). Previously, Impala would attempt to read the file despite the possibility that the file was corrupted.
- Previously, calls to overloaded built-in functions could treat parameters as `DOUBLE` or `FLOAT` when no overload had a signature that matched the exact argument types. Now Impala prefers the function signature with `DECIMAL` parameters in this case. This change avoids a possible loss of precision in function calls such as `greatest(0, 99999.8888)`; now both parameters are treated as `DECIMAL` rather than `DOUBLE`, avoiding any loss of precision in the fractional value. This could cause slightly different results than in previous Impala releases for certain function calls.
- Formerly, adding or subtracting a large interval value to a `TIMESTAMP` could produce a nonsensical result. Now when the result goes outside the range of `TIMESTAMP` values, Impala returns `NULL`.
- Formerly, it was possible to accidentally create a table with identical row and column delimiters. This could happen unintentionally, when specifying one of the delimiters and using the default value for the other. Now an attempt to use identical delimiters still succeeds, but displays a warning message.
- Formerly, Impala could include snippets of table data in log files by default, for example when reporting conversion errors for data values. Now any such log messages are only produced at higher logging levels that you would enable only during debugging.

Impala Incompatible Changes Introduced in CDH 5.4.x



Note: Impala 2.2.0 is available as part of CDH 5.4.0 and is not available for CDH 4. Cloudera does not intend to release future versions of Impala for CDH 4 outside patch and maintenance releases if required. Given the end-of-maintenance status for CDH 4, Cloudera recommends all customers to migrate to a recent CDH 5 release.

Changes to File Handling

Impala queries ignore files with extensions commonly used for temporary work files by Hadoop tools. Any files with extensions `.tmp` or `.copying` are not considered part of the Impala table. The suffix matching is case-insensitive, so for example Impala ignores both `.copying` and `.COPYING` suffixes.

The log rotation feature in Impala 2.2.0 and higher means that older log files are now removed by default. The default is to preserve the latest 10 log files for each severity level, for each Impala-related daemon. If you have set up your own log rotation processes that expect older files to be present, either adjust your procedures or change the Impala `-max_log_files` setting.

Changes to Prerequisites

The prerequisite for CPU architecture has been relaxed in Impala 2.2.0 and higher. From this release onward, Impala works on CPUs that have the SSE3 instruction set. The SSE4 instruction set is no longer required. This relaxed requirement simplifies the upgrade planning from Impala 1.x releases, which also worked on SSE3-enabled processors.

Incompatible Changes Introduced in CDH 5.3.x

Changes to Prerequisites

Currently, Impala 2.1.x does not function on CPUs without the SSE4.1 instruction set. This minimum CPU requirement is higher than in previous versions, which relied on the older SSE3 instruction set. Check the CPU level of the hosts in your cluster before upgrading to CDH 5.3 / Impala 2.1.

Changes to Output Format

The “small query” optimization feature introduces some new information in the `EXPLAIN` plan, which you might need to account for if you parse the text of the plan output.

New Reserved Words

New SQL syntax introduces additional reserved words: `FOR`, `GRANT`, `REVOKE`, `ROLE`, `ROLES`, `INCREMENTAL`.

Incompatible Changes Introduced in Impala 2.0.5 / CDH 5.2.6

No incompatible changes.



Note: Impala 2.0.5 is available as part of CDH 5.2.6, not under CDH 4.

Incompatible Changes Introduced in Impala 2.0.4 / CDH 5.2.5

No incompatible changes.



Note: Impala 2.0.4 is available as part of CDH 5.2.5, not under CDH 4.

Incompatible Changes Introduced in Impala 2.0.3 / CDH 5.2.4



Note: Impala 2.0.3 is available as part of CDH 5.2.4, not under CDH 4.

Incompatible Changes Introduced in Impala 2.0.2 / CDH 5.2.3

No incompatible changes.



Note: Impala 2.0.2 is available as part of CDH 5.2.3, not under CDH 4.

Incompatible Changes Introduced in Impala 2.0.1 / CDH 5.2.1

- The `INSERT` statement has always left behind a hidden work directory inside the data directory of the table. Formerly, this hidden work directory was named `.impala_insert_staging`. In Impala 2.0.1 and later, this directory name is changed to `_impala_insert_staging`. (While HDFS tools are expected to treat names beginning either with underscore and dot as hidden, in practice names beginning with an underscore are more widely supported.) If you have any scripts, cleanup jobs, and so on that rely on the name of this work directory, adjust them to use the new name.
- The `abs()` function now takes a broader range of numeric types as arguments, and the return type is the same as the argument type.
- Shorthand notation for character classes in regular expressions, such as `\d` for digit, are now available again in regular expression operators and functions such as `regexp_extract()` and `regexp_replace()`. Some other differences in regular expression behavior remain between Impala 1.x and Impala 2.x releases. See [Incompatible Changes Introduced in Impala 2.0.0 / CDH 5.2.0](#) on page 1931 for details.

Incompatible Changes Introduced in Impala 2.0.0 / CDH 5.2.0

Changes to Prerequisites

Currently, Impala 2.0.x does not function on CPUs without the SSE4.1 instruction set. This minimum CPU requirement is higher than in previous versions, which relied on the older SSSE3 instruction set. Check the CPU level of the hosts in your cluster before upgrading to CDH 5.2 / Impala 2.0.

Changes to Query Syntax

The new syntax where query hints are allowed in comments causes some changes in the way comments are parsed in the `impala-shell` interpreter. Previously, you could end a `--` comment line with a semicolon and `impala-shell` would treat that as a no-op statement. Now, a comment line ending with a semicolon is passed as an empty statement to the Impala daemon, where it is flagged as an error.

Impala 2.0 and later uses a different support library for regular expression parsing than in earlier Impala versions. Now, Impala uses the [Google RE2 library](#) rather than Boost for evaluating regular expressions. This implementation change causes some differences in the allowed regular expression syntax, and in the way certain regex operators are interpreted. The following are some of the major differences (not necessarily a complete list):

- `. * ?` notation for non-greedy matches is now supported, where it was not in earlier Impala releases.
- By default, `^` and `$` now match only begin/end of buffer, not begin/end of each line. This behavior can be overridden in the regex itself using the `m` flag.
- By default, `.` does not match newline. This behavior can be overridden in the regex itself using the `s` flag.
- `\z` is not supported.

- < and > for start of word and end of word are not supported.
- Lookahead and lookbehind are not supported.
- Shorthand notation for character classes, such as \d for digit, is not recognized. (This restriction is lifted in Impala 2.0.1, which restores the shorthand notation.)

Changes to Output Format

In Impala 2.0 and later, `user ()` returns the full Kerberos principal string, such as `user@example.com`, in a Kerberized environment.

The changed format for the user name in secure environments is also reflected where the user name is displayed in the output of the `PROFILE` command.

In the output from `SHOW FUNCTIONS`, `SHOW AGGREGATE FUNCTIONS`, and `SHOW ANALYTIC FUNCTIONS`, arguments and return types of arbitrary `DECIMAL` scale and precision are represented as `DECIMAL(*, *)`. Formerly, these items were displayed as `DECIMAL(-1, -1)`.

Changes to Query Options

The `PARQUET_COMPRESSION_CODEC` query option has been replaced by the `COMPRESSION_CODEC` query option.

Changes to Configuration Options

The meaning of the `--idle_query_timeout` configuration option is changed, to accommodate the new `QUERY_TIMEOUT_S` query option. Rather than setting an absolute timeout period that applies to all queries, it now sets a maximum timeout period, which can be adjusted downward for individual queries by specifying a value for the `QUERY_TIMEOUT_S` query option. In sessions where no `QUERY_TIMEOUT_S` query option is specified, the `--idle_query_timeout` timeout period applies the same as in earlier versions.

The `--strict_unicode` option of `impala-shell` was removed. To avoid problems with Unicode values in `impala-shell`, define the following locale setting before running `impala-shell`:

```
export LC_CTYPE=en_US.UTF-8
```

New Reserved Words

Some new SQL syntax requires the addition of new reserved words: `ANTI`, `ANALYTIC`, `OVER`, `PRECEDING`, `UNBOUNDED`, `FOLLOWING`, `CURRENT`, `ROWS`, `RANGE`, `CHAR`, `VARCHAR`.

Changes to Data Files

The default Parquet block size for Impala is changed from 1 GB to 256 MB. This change could have implications for the sizes of Parquet files produced by `INSERT` and `CREATE TABLE AS SELECT` statements.

Although older Impala releases typically produced files that were smaller than the old default size of 1 GB, now the file size matches more closely whatever value is specified for the `PARQUET_FILE_SIZE` query option. Thus, if you use a non-default value for this setting, the output files could be larger than before. They still might be somewhat smaller than the specified value, because Impala makes conservative estimates about the space needed to represent each column as it encodes the data.

When you do not specify an explicit value for the `PARQUET_FILE_SIZE` query option, Impala tries to keep the file size within the 256 MB default size, but Impala might adjust the file size to be somewhat larger if needed to accommodate the layout for *wide* tables, that is, tables with hundreds or thousands of columns.

This change is unlikely to affect memory usage while writing Parquet files, because Impala does not pre-allocate the memory needed to hold the entire Parquet block.

Incompatible Changes Introduced in Impala 1.4.4 / CDH 5.1.5

No incompatible changes.



Note: Impala 1.4.4 is available as part of CDH 5.1.5, not under CDH 4.

Incompatible Changes Introduced in Impala 1.4.3 / CDH 5.1.4

No incompatible changes. The TLS/SSL security fix does not require any change in the way you interact with Impala.



Note: Impala 1.4.3 is available as part of CDH 5.1.4, and under CDH 4.

Incompatible Changes Introduced in Impala 1.4.2 / CDH 5.1.3

None. Impala 1.4.2 is purely a bug-fix release. It does not include any incompatible changes.



Note: Impala 1.4.2 is only available as part of CDH 5.1.3, not under CDH 4.

Incompatible Changes Introduced in Impala 1.4.1 / CDH 5.1.2

None. Impala 1.4.1 is purely a bug-fix release. It does not include any incompatible changes.

Incompatible Changes Introduced in Impala 1.4.0 / CDH 5.1.0

- There is a slight change to required security privileges in the Sentry framework. To create a new object, now you need the `ALL` privilege on the parent object. For example, to create a new table, view, or function requires having the `ALL` privilege on the database containing the new object.
- With the ability of `ORDER BY` queries to process unlimited amounts of data with no `LIMIT` clause, the query options `DEFAULT_ORDER_BY_LIMIT` and `ABORT_ON_DEFAULT_LIMIT_EXCEEDED` are now deprecated and have no effect.
- There are some changes to the list of reserved words. The following keywords are new:
 - `API_VERSION`
 - `BINARY`
 - `CACHED`
 - `CLASS`
 - `PARTITIONS`
 - `PRODUCED`
 - `UNCACHED`

The following were formerly reserved keywords, but are no longer reserved:

- `COUNT`
 - `GROUP_CONCAT`
 - `NDV`
 - `SUM`
- The fix for issue [IMPALA-973](#) changes the behavior of the `INVALIDATE METADATA` statement regarding nonexistent tables. In Impala 1.4.0 and higher, the statement returns an error if the specified table is not in the metastore database at all. It completes successfully if the specified table is in the metastore database but not yet recognized by Impala, for example if the table was created through Hive. Formerly, you could issue this statement for a completely nonexistent table, with no error.

Incompatible Changes Introduced in Impala 1.3.3 / CDH 5.0.5

No incompatible changes. The TLS/SSL security fix does not require any change in the way you interact with Impala.



Note: Impala 1.3.3 is only available as part of CDH 5.0.5, not under CDH 4.

Incompatible Changes Introduced in Impala 1.3.2 / CDH 5.0.4

With the fix for IMPALA-1019, you can use HDFS caching for files that are accessed by Impala.



Note: Impala 1.3.2 is only available as part of CDH 5.0.4, not under CDH 4.

Incompatible Changes Introduced in Impala 1.3.1 / CDH 5.0.3

- In Impala 1.3.1 and higher, the `REGEXP` and `RLIKE` operators now match a regular expression string that occurs anywhere inside the target string, the same as if the regular expression was enclosed on each side by `.*`. See [REGEXP Operator](#) for examples. Previously, these operators only succeeded when the regular expression matched the entire target string. This change improves compatibility with the regular expression support for popular database systems. There is no change to the behavior of the `regexp_extract()` and `regexp_replace()` built-in functions.
- The result set for the `SHOW FUNCTIONS` statement includes a new first column, with the data type of the return value.

Incompatible Changes Introduced in Impala 1.3.0 / CDH 5.0.0

- The `EXPLAIN_LEVEL` query option now accepts numeric options from 0 (most concise) to 3 (most verbose), rather than only 0 or 1. If you formerly used `SET EXPLAIN_LEVEL=1` to get detailed explain plans, switch to `SET EXPLAIN_LEVEL=3`. If you used the mnemonic keyword (`SET EXPLAIN_LEVEL=verbose`), you do not need to change your code because now level 3 corresponds to `verbose`.
- The keyword `DECIMAL` is now a reserved word. If you have any databases, tables, columns, or other objects already named `DECIMAL`, quote any references to them using backticks (```) to avoid name conflicts with the keyword.



Note: Although the `DECIMAL` keyword is a reserved word, currently Impala does not support `DECIMAL` as a data type for columns.

- The query option formerly named `YARN_POOL` is now named `REQUEST_POOL` to reflect its broader use with the Impala admission control feature.
- There are some changes to the list of reserved words.
 - The names of aggregate functions are no longer reserved words, so you can have databases, tables, columns, or other objects named `AVG`, `MIN`, and so on without any name conflicts.
 - The internal function names `DISTINCTPC` and `DISTINCTPCSA` are no longer reserved words, although `DISTINCT` is still a reserved word.
 - The keywords `CLOSE_FN` and `PREPARE_FN` are now reserved words.
- The HDFS property `dfs.client.file-block-storage-locations.timeout` was renamed to `dfs.client.file-block-storage-locations.timeout.millis`, to emphasize that the unit of measure is milliseconds, not seconds. Impala requires a timeout of at least 10 seconds, making the minimum value for this setting 10000. On systems not managed by Cloudera Manager, you might need to edit the `hdfs-site.xml` file in the Impala configuration directory for the new name and minimum value.

Incompatible Changes Introduced in Impala 1.2.4

There are no incompatible changes introduced in Impala 1.2.4.

Previously, after creating a table in Hive, you had to issue the `INVALIDATE METADATA` statement with no table name, a potentially expensive operation on clusters with many databases, tables, and partitions. Starting in Impala 1.2.4, you can issue the statement `INVALIDATE METADATA table_name` for a table newly created through Hive. Loading the metadata for only this one table is faster and involves less network overhead. Therefore, you might revisit your setup DDL scripts to add the table name to `INVALIDATE METADATA` statements, in cases where you create and populate the tables through Hive before querying them through Impala.

Incompatible Changes Introduced in Impala 1.2.3

Because the feature set of Impala 1.2.3 is identical to Impala 1.2.2, there are no new incompatible changes. See [Incompatible Changes Introduced in Impala 1.2.2](#) on page 1935 if you are upgrading from Impala 1.2.1 or 1.1.x.

Incompatible Changes Introduced in Impala 1.2.2

The following changes to SQL syntax and semantics in Impala 1.2.2 could require updates to your SQL code, or schema objects such as tables or views:

- With the addition of the `CROSS JOIN` keyword, you might need to rewrite any queries that refer to a table named `CROSS` or use the name `CROSS` as a table alias:

```
-- Formerly, 'cross' in this query was an alias for t1
-- and it was a normal join query.
-- In 1.2.2 and higher, CROSS JOIN is a keyword, so 'cross'
-- is not interpreted as a table alias, and the query
-- uses the special CROSS JOIN processing rather than a
-- regular join.
select * from t1 cross join t2...

-- Now if CROSS is used in other context such as a table or column name,
-- use backticks to escape it.
create table `cross` (x int);
select * from `cross`;
```

- Formerly, a `DROP DATABASE` statement in Impala would not remove the top-level HDFS directory for that database. The `DROP DATABASE` has been enhanced to remove that directory. (You still need to drop all the tables inside the database first; this change only applies to the top-level directory for the entire database.)
- The keyword `PARQUET` is introduced as a synonym for `PARQUETFILE` in the `CREATE TABLE` and `ALTER TABLE` statements, because that is the common name for the file format. (As opposed to `SequenceFile` and `RCFile` where the “File” suffix is part of the name.) Documentation examples have been changed to prefer the new shorter keyword. The `PARQUETFILE` keyword is still available for backward compatibility with older Impala versions.
- New overloads are available for several operators and built-in functions, allowing you to insert their result values into smaller numeric columns such as `INT`, `SMALLINT`, `TINYINT`, and `FLOAT` without using a `CAST()` call. If you remove the `CAST()` calls from `INSERT` statements, those statements might not work with earlier versions of Impala.

Because many users are likely to upgrade straight from Impala 1.x to Impala 1.2.2, also read [Incompatible Changes Introduced in Impala 1.2.1](#) on page 1935 for things to note about upgrading to Impala 1.2.x in general.

Incompatible Changes Introduced in Impala 1.2.1

The following changes to SQL syntax and semantics in Impala 1.2.1 could require updates to your SQL code, or schema objects such as tables or views:

- In Impala 1.2.1 and higher, all `NULL` values come at the end of the result set for `ORDER BY ... ASC` queries, and at the beginning of the result set for `ORDER BY ... DESC` queries. In effect, `NULL` is considered greater than all other values for sorting purposes. The original Impala behavior always put `NULL` values at the end, even for `ORDER BY ... DESC` queries. The new behavior in Impala 1.2.1 makes Impala more compatible with other popular database systems. In Impala 1.2.1 and higher, you can override or specify the sorting behavior for `NULL` by adding the clause `NULLS FIRST` or `NULLS LAST` at the end of the `ORDER BY` clause.

The new `catalogd` service might require changes to any user-written scripts that stop, start, or restart Impala services, install or upgrade Impala packages, or issue `REFRESH` or `INVALIDATE METADATA` statements:

- See [Setting Up Apache Impala Using the Command Line](#), [Upgrading Impala](#) and [Starting Impala](#), for usage information for the `catalogd` daemon.
- The `REFRESH` and `INVALIDATE METADATA` statements are no longer needed when the `CREATE TABLE`, `INSERT`, or other table-changing or data-changing operation is performed through Impala. These statements are still needed if such operations are done through Hive or by manipulating data files directly in HDFS, but in those cases the statements only need to be issued on one Impala node rather than on all nodes. See [REFRESH Statement](#) and [INVALIDATE METADATA Statement](#) for the latest usage information for those statements.
- See [The Impala Catalog Service](#) for background information on the `catalogd` service.

Incompatible Changes Introduced in Impala 1.2.0 (Beta)

There are no incompatible changes to SQL syntax in Impala 1.2.0 (beta).

Because Impala 1.2.0 is bundled with the CDH 5 beta download and depends on specific levels of Apache Hadoop components supplied with CDH 5, you can only install it in combination with the CDH 5 beta.

The new `catalogd` service might require changes to any user-written scripts that stop, start, or restart Impala services, install or upgrade Impala packages, or issue `REFRESH` or `INVALIDATE METADATA` statements:

- See [Setting Up Apache Impala Using the Command Line](#), [Upgrading Impala](#) and [Starting Impala](#), for usage information for the `catalogd` daemon.
- The `REFRESH` and `INVALIDATE METADATA` statements are no longer needed when the `CREATE TABLE`, `INSERT`, or other table-changing or data-changing operation is performed through Impala. These statements are still needed if such operations are done through Hive or by manipulating data files directly in HDFS, but in those cases the statements only need to be issued on one Impala node rather than on all nodes. See [REFRESH Statement](#) and [INVALIDATE METADATA Statement](#) for the latest usage information for those statements.
- See [The Impala Catalog Service](#) for background information on the `catalogd` service.

The new resource management feature interacts with both YARN and Llama services, which are available in CDH 5. These services are set up for you automatically in a Cloudera Manager (CM) environment. For information about setting up the YARN and Llama services, see the instructions for [YARN](#) and [Llama](#) in the *CDH 5 Documentation*.

Incompatible Changes Introduced in Impala 1.1.1

There are no incompatible changes in Impala 1.1.1.

Previously, it was not possible to create Parquet data through Impala and reuse that table within Hive. Now that Parquet support is available for Hive 10, reusing existing Impala Parquet data files in Hive requires updating the table metadata. Use the following command if you are already running Impala 1.1.1:

```
ALTER TABLE table_name SET FILEFORMAT PARQUETFILE;
```

If you are running a level of Impala that is older than 1.1.1, do the metadata update through Hive:

```
ALTER TABLE table_name SET SERDE 'parquet.hive.serde.ParquetHiveSerDe';
ALTER TABLE table_name SET FILEFORMAT
  INPUTFORMAT "parquet.hive.DeprecatedParquetInputFormat"
  OUTPUTFORMAT "parquet.hive.DeprecatedParquetOutputFormat";
```

Impala 1.1.1 and higher can reuse Parquet data files created by Hive, without any action required.

As usual, make sure to upgrade the Impala LZO support package to the latest level at the same time as you upgrade the Impala server.

Incompatible Change Introduced in Impala 1.1

- The `REFRESH` statement now requires a table name; in Impala 1.0, the table name was optional. This syntax change is part of the internal rework to make `REFRESH` a true Impala SQL statement so that it can be called through the JDBC and ODBC APIs. `REFRESH` now reloads the metadata immediately, rather than marking it for update the next time any affected table is accessed. The previous behavior, where omitting the table name caused a refresh of the entire Impala metadata catalog, is available through the new `INVALIDATE METADATA` statement. `INVALIDATE METADATA` can be specified with a table name to affect a single table, or without a table name to affect the entire metadata catalog; the relevant metadata is reloaded the next time it is requested during the processing for a SQL statement. See [REFRESH Statement](#) and [INVALIDATE METADATA Statement](#) for the latest details about these statements.

Incompatible Changes Introduced in Impala 1.0

- If you use LZO-compressed text files, when you upgrade Impala to version 1.0, also update the `impala-lzo-cdh4` to the latest level. See [Using LZO-Compressed Text Files](#) for details.
- Cloudera Manager 4.5.2 and higher only supports Impala 1.0 and higher, and vice versa. If you upgrade to Impala 1.0 or higher managed by Cloudera Manager, you must also upgrade Cloudera Manager to version 4.5.2 or higher. If you upgrade from an earlier version of Cloudera Manager, and were using Impala, you must also upgrade Impala to version 1.0 or higher. The beta versions of Impala are no longer supported as of the release of Impala 1.0.

Cloudera Distribution of Apache Kafka Incompatible Changes and Limitations

Flume shipped with CDH 5.6.0 can only send data to Kafka 2.0 and higher via unsecured transport

To take advantage of security additions in Kafka 2.0 and higher, upgrade to CDH 5.7.0.

Topic Blacklist Removed

The MirrorMaker **Topic blacklist** setting has been removed in Cloudera Distribution of Kafka 2.x and higher.

Avoid Data Loss Option Removed

The **Avoid Data Loss** option from earlier releases has been removed in Kafka 2.x in favor of automatically setting the following properties.

1. Producer settings
 - `acks=all`
 - `retries=max integer`
 - `max.block.ms=max long`
2. Consumer setting
 - `auto.commit.enable=false`
3. MirrorMaker setting
 - `abort.on.send.failute=true`

Kite Incompatible Changes and Limitations

Kite in CDH has been rebased on the 1.0 release upstream. This breaks backward compatibility with existing APIs. The APIs are documented at <http://kitesdk.org/docs/1.0.0/apidocs/index.html>. For more information, see [What's New In CDH 5.4.x](#) on page 1838.

Apache Kudu Incompatible Changes and Limitations

See the Apache Kudu [New Features and Changes](#) topic for any incompatible changes introduced in each release.

For a complete list of limitations in the latest release of Apache Kudu, see [Apache Kudu Usage Limitations](#).

For limitations in older versions (Kudu 1.4.0 and lower), see the documentation for previous versions of [Apache Kudu \(1.4.0 and lower\)](#).

Llama Incompatible Changes and Limitations



Note: Llama no longer supported:

Although Impala can be used together with YARN via simple configuration of Static Service Pools in Cloudera Manager, the use of the general-purpose component Llama for integrated resource management between Impala and YARN is no longer supported as of CDH 5.5 / Impala 2.3. Please contact your Cloudera account team if this impacts a previously deployed system.

The following changes have made in the Llama API to help solve synchronization problems:

- As of CDH 5.1, the Reserve API requires you to provide the `reservationId` as part of the request. In previous releases, the `reservationId` was auto-generated and returned to the user.

The `TLLamaAMReservationRequest` has an additional field called `reservation_id` which needs to be initialized to a UUID value. If you do not set this field, the request will result in an error with the error code set to `ErrorCode.RESERVATION_NO_ID_PROVIDED`.

- The Expand API now requires you to provide the `expansionId` as part of the request. In previous releases, the `expansionId` was auto-generated and returned to the user.

The `TLLamaAMReservationExpansionRequest` has an additional field called `expansion_id` which needs to be initialized to a UUID value. If you do not set this field, the request will result in an error with the error code set to `ErrorCode.EXPANSION_NO_EXPANSION_ID_PROVIDED`

Apache Mahout Incompatible Changes and Limitations

CDH 5.1 introduces the following incompatible changes.

Minor changes in behavior:

- [MAHOUT-1368](#)

The `org.apache.mahout.math.stats.OnlineSummarizer` algorithm has changed, potentially leading to different results.

- [MAHOUT-1392](#)

The streaming `KMeans` implementation has changed in how it outputs centroids when run outside of a Hadoop cluster.

- [MAHOUT-1565](#)

Major Developer API changes

- [MAHOUT-1296](#):

The following implementations have been removed:

- Clustering:
 - `DirichletMeanShift`
 - `MinHash`
 - `Eigencuts in o.a.m.clustering.spectral.eigencuts`
- Classification:
 - `WinnowPerceptron`
 - `Frequent Pattern Mining`

- Collaborative Filtering:
 - All recommenders in `o.a.m.cf.taste.impl.recommender.knn`
 - `TreeClusteringRecommender` in `o.a.m.cf.taste.impl.recommender`
 - SlopeOne implementations in `o.a.m.cf.taste.hadoop.slopeone` and `o.a.m.cf.taste.impl.recommender.slopeone`
 - Distributed pseudo recommender in `o.a.m.cf.taste.hadoop.pseudo`
- [MAHOUT-1362](#):
The `examples/bin/build-reuters.sh` script has been removed.

Minor Developer API changes

- [MAHOUT-1280](#):
Some SSVD support code, such as `UpperTriangularMatrix`, has been moved to `mahout-math`.
- [MAHOUT-1363](#):
Scala-related math code has been moved into an `org.apache.mahout.math.scalabindings` sub-package.

Minor packaging changes

- [MAHOUT-1382](#):
Move to Guava 16.
- [MAHOUT-1364](#):
Update to require Lucene 4.6.1.

Apache Oozie Incompatible Changes and Limitations

As of CDH 5.2.0, a new [Hive 2 Action](#) allows Oozie to run HiveServer2 scripts. Using the Hive Action with HiveServer2 is now deprecated; you should switch to the new Hive 2 Action as soon as possible.

CDH 5.4.0 introduces sharelib packaging changes: the sharelib was previously shipped as a pair of tarballs, `oozie-sharelib-yarn.tar.gz` and `oozie-sharelib-mr1.tar.gz`. As of CDH 5.4.0, it is shipped as a pair of directories, `oozie-sharelib-yarn` and `oozie-sharelib-mr1`. They are still installed in `/usr/lib/oozie/` in a packages distribution, and in `/lib/oozie` in a parcels distribution.

Apache Pig Incompatible Changes and Limitations

- Apache Pig has been upgraded from version 0.11 to 0.12.
- A custom UDF must return the schema as a tuple with exactly one field.
- Added the `In`, `CASE` and `Assert` keywords; They cannot be used as variables or UDF names any more.

Cloudera Search Incompatible Changes and Limitations

General Limitations of Cloudera Search

- Cloudera Search supports one instance of the Apache Solr service on each host in a cluster. Using multiple Solr instances on a host is not supported.
- If multiple Solr instances are configured to depend on the same Sentry service, it is not possible to create unique Solr Sentry privileges per Solr deployment. Since privileges are enforced in all Solr instances simultaneously, you cannot add distinct privileges that apply to one Solr cluster, but not to another.
- **Converting existing file-based Sentry authorization policy files to permissions in the Sentry service does not support preserving case-sensitive role or group names**

The file-based model allows case-sensitive role names. During conversion, all roles and groups are converted to lower case.

- If a policy-file conversion will change the case of roles or groups, a warning is presented. Policy conversion can proceed, but if you have enabled document-level security and use role names as your tokens, you must re-index using the new lower case role names after conversion is complete.
- If a policy-file conversion will change the case of roles or groups, creating a name collision, an error occurs and conversion cannot occur. In such a case, you must eliminate the collisions before proceeding. For example, you could rename or delete the roles or groups that cause a collision.

Incompatible Changes Between Cloudera Search for CDH 5.8 and Previous Versions of Cloudera Search

- **Converting existing file-based Sentry authorization policy files to permissions in the Sentry service does not support preserving case-sensitive role or group names**

The file-based model allows case-sensitive role names. During conversion, all roles and groups are converted to lower case.

- If a policy-file conversion will change the case of roles or groups, a warning is presented. Policy conversion can proceed, but if you have enabled document-level security and use role names as your tokens, you must re-index using the new lower case role names after conversion is complete.
- If a policy-file conversion will change the case of roles or groups, creating a name collision, an error occurs and conversion cannot occur. In such a case, you must eliminate the collisions before proceeding. For example, you could rename or delete the roles or groups that cause a collision.

Incompatible Changes Between Cloudera Search for CDH 5.5 and Previous Versions of Cloudera Search

- **Using MapReduceIndexerTool with configurations that require an `updateRequestProcessorChain` may fail unless an alternate configuration is specified**

With Search for CDH 5.5, the MapReduceIndexerTool uses a default `solrconfig.xml` that is appropriate for most collection configurations. With this configuration, the MapReduceIndexerTool can index data even if Sentry is enabled. This default configuration does not include any `updateRequestProcessorChains`; if your configuration requires an `updateRequestProcessorChain`, you can configure the MapReduceIndexerTool to use the configuration from ZooKeeper by specifying `--use-zk-solrconfig.xml` or from local disk by specifying `--solr-home-dir`.

Incompatible Changes Between Cloudera Search for CDH 5.4 and Previous Versions of Cloudera Search

- **HDFS locality metrics are disabled by default**

HDFS locality metrics are disabled by default in CDH 5.4.8 and higher because they can generate numerous HDFS calls, negatively affecting performance. Performance degradation is more common in production-scale environments that include rapidly changing indexes. To re-enable these metrics, add the following to the directory config in `solrconfig.xml`:

```
<directoryFactory name="DirectoryFactory"
class="${solr.directoryFactory:org.apache.solr.core.HdfsDirectoryFactory}">
  <bool name="solr.hdfs.locality.metrics.enabled">true</bool>
</directoryFactory>
```

- **CloudSolrServer and LBHttpSolrServer no longer declare `MalformedURLException` as thrown from their constructors**

As a result of this change, compilation failures against the 4.10.3 Solr libraries may fail. To avoid this issue, make relevant source code changes, such as removing catch phrases related to `MalformedURLException`, and then recompile the application.

Related JIRA: Solr-5555

- **The `solrJ` client `JavaBinCodec` serializes unknown objects differently**

Starting with Search for CDH 5.4.0, Search moves from Solr 4.4 to Solr 4.1.0. With Solr 4.4, JavaBinCodec serialized unknown Java objects as `obj.toString()`. In Solr 4.10.0, JavaBinCodec serializes unknown Java objects as `obj.getClass().getName() + ':' + obj.toString()`.

As a result, the same objects may produce different results when serialized with CDH 5.4 and higher compared with objects serialized with CDH 5.3 and lower.

- **Parsing using `schema.xml` creates an init error when `<dynamicField/>` declarations include `default` or `required` attributes**

In previous releases, these attributes were ignored. If init errors occur when upgrading with an existing `schema.xml`, remove the `default` or `required` attributes. After removing these attributes, Search functions as it did before upgrading.

Related JIRA: SOLR-5227.

- **Indexing documents with terms that exceed Lucene's `MAX_TERM_LENGTH` registers errors**

In previous releases, terms that exceeded the length limit were silently ignored. To make Search function as it did in previous releases—silently ignoring longer terms—use `solr.LengthFilterFactory` in all of your Analyzers.

Related JIRA: LUCENE-5472.

- **The `fieldType` configuration `docValuesFormat="Disk"` is no longer supported**

If your `schema.xml` contains `fieldTypes` using `docValuesFormat="Disk"`, modify the file to remove the `docValuesFormat` attribute and optimize your index to rewrite to the default codec. Make these changes before upgrading to CDH 5.4.

Related JIRA: LUCENE-5761.

- **`UpdateRequestExt` has been removed**

Use `UpdateRequest` instead.

Related JIRA: SOLR-4816.

- **Parsing `schema.xml` registers errors when multiple values exist where only a single value is permitted.**

With previous releases, when multiple values existed where only a single value was permitted, one value was silently chosen. In CDH 5.4, if multiple values exist where only a single value is supported, configuration parsing fails. The extra values must be removed.

Related JIRAs: SOLR-4953, SOLR-5108.

Incompatible Changes Between Cloudera Search for CDH 5.2 and Cloudera Search for CDH 5.3

Some packaging changes were made that have consequences for `CrunchIndexerTool` start-up scripts. If those startup scripts include the following line:

```
export myDriverJar=$(find $myDriverJarDir -maxdepth 1 -name \
'*.*jar' ! -name '*-job.jar' ! -name '*-sources.jar')
```

That line in those scripts should be changed as follows:

```
export myDriverJar=$(find $myDriverJarDir -maxdepth 1 -name \
'search-crunch-*.jar' ! -name '*-job.jar' ! -name '*-sources.jar')
```

Incompatible changes between Cloudera Search for CDH 5 Beta 2 and Older Versions of Cloudera Search

The following incompatible changes occurred between Cloudera Search for CDH 5 beta 2 and older versions of Cloudera Search, including both lower versions of Cloudera Search for CDH 5 and Cloudera Search 1.x:

- Supported values for the `--reducers` option of the `MapReduceIndexer` tool change with the release of Search for CDH 5 beta 2. To use one reducer per output shard, 0 is used in Search 1.x and Search for CDH 5 beta 1. With the release of Search for CDH 5 beta 2, -2 is used for one reducer per output shard. Because of this change,

commands using `--reducers 0` that were written for previous Search releases do not continue to work in the same way after upgrading to Search for CDH 5 beta 2. After upgrading to Search for CDH 5 beta 2, using `--reducers 0` results in an exception stating that zero is an illegal value.

Apache Sentry Incompatible Changes

The compatibility notes for Sentry are listed below, organized by CDH version.

Compatibility Notes for CDH 5.13.x

CDH 5.13.0 introduces high availability for Sentry. Although you can still use Sentry with a flat file database, if you want to configure high availability for Sentry, the Sentry service must use by a relational database.

Compatibility Notes for CDH 5.3.x

As of CDH 5.3, `MSCK REPAIR TABLE` now requires `ALL` privileges for the table (previously this statement required `ALL` privileges on the parent database for the table).

Compatibility Notes for CDH 5.2.x

Upgrading Sentry from a release earlier than CHD 5.2 to CDH 5.2 or later entails a schema upgrade to the Sentry database.

Compatibility Notes for CDH 5.1.x

CDH 5.1 introduces a new privilege model in Sentry. This introduces a backward incompatible change for Impala. Creating a new object now requires the `ALL` privilege on the parent object. For example, creating a database now requires server-level privileges (previously needed database-level) and creating a table requires database-level privileges (previously needed table-level).

Apache Spark Incompatible Changes and Limitations

- If you have uploaded the Spark assembly JAR file to HDFS, you must upload the new version of the file each time you upgrade Spark to a new *minor* CDH release (for example, any CDH 5.2, 5.3, 5.4, or 5.5 release). You may also need to modify the configured path for the file; see the instructions under the CDH 5.2 section.
- **CDH 5.5**
 - Dynamic allocation is enabled by default but is not compatible with streaming. For streaming jobs, if you specify `--num-executors`, then dynamic allocation is implicitly disabled. To be safe, you can explicitly disable dynamic allocation using: `spark.dynamicAllocation.enabled = false`
 - The CDH 5.5 version of Spark 1.5 differs from the Apache Spark 1.5 release in using Akka version 2.2.3, the version used by Spark 1.1 and CDH 5.2. Apache Spark 1.5 uses Akka version 2.3.11.
- **CDH 5.4**
 - The CDH 5.4 version of Spark 1.3 differs from the Apache Spark 1.3 release in using Akka version 2.2.3, the version used by Spark 1.1 and CDH 5.2. Apache Spark 1.3 uses Akka version 2.3.4.
- **CDH 5.3**
 - Spark 1.2, on which CDH 5.3 is based, does not expose a transitive dependency on the Guava library. As a result, projects that use Guava but do not explicitly add it as a dependency will need to be modified: the dependency must be added to the project and also packaged with the job.
 - The CDH 5.3 version of Spark 1.2 differs from the Apache Spark 1.2 release in using Akka version 2.2.3, the version used by Spark 1.1 and CDH 5.2. Apache Spark 1.2 uses Akka version 2.3.4.
- **CDH 5.2**
 - The configured paths for `spark.eventLog.dir`, `spark.history.fs.logDirectory`, and the `SPARK_JAR` environment variable have changed in a way that may not be backward-compatible. By default, those paths now refer to the local filesystem. To make sure everything works as before, modify the paths as follows:
 - For HDFS, if this is not a federated cluster, prepend `hdfs:` to the path.

- For HDFS in a federated cluster, prepend `viewfs:` to the path.

Alternatively, you can prepend the value of `fs.defaultFS`, set in `core-site.xml` in the HDFS configuration.

- The following changes may affect existing applications:
 - The default for I/O compression is now Snappy (changed from LZF).
 - PySpark now performs external spilling during aggregations.
- The following Spark-related artifacts are no longer published as part of the Cloudera repository:
 - `spark-assembly`: The `spark-assembly` JAR is used internally by Spark distributions when running Spark applications and should not be referenced directly. Instead, projects should add dependencies for those parts of the Spark project that are being used, for example, `spark-core`.
 - `spark-yarn`
 - `spark-tools`
 - `spark-examples`
 - `spark-repl`

• CDH 5.1

- Before you can run Spark in standalone mode, you must set the `spark.master` property in `/etc/spark/conf/spark-defaults.conf`, as follows:

```
spark.master=spark://MASTER_IP:MASTER_PORT
```

where `MASTER_IP` is the IP address of the host the Spark master is running on and `MASTER_PORT` is the port.

This setting means that all jobs will run in standalone mode by default; you can override the default on the command line.

- Includes changes that will enable Spark to avoid breaking compatibility in the future. As a result, most applications will require a recompile to run against Spark 1.0, and some will require changes in source code. The details are as follows:
 - This release has two changes in the core Scala API:
 - The `cogroup` and `groupByKey` operators now return an `Iterator` over their values instead of a `Seq`. This change means that the set of values corresponding to a particular key need not all reside in memory at the same time.
 - `SparkContext.jarOfClass` now returns `Option[String]` instead of `Seq[String]`.
 - Spark Java APIs have been updated to accommodate Java 8 lambdas. See [Migrating from pre-1.0 Versions of Spark](#) for more information.

Apache Sqoop Incompatible Changes and Limitations

- The Sqoop `import-mainframe` feature is not supported in CDH at all.
- Sqoop 2 is deprecated. Cloudera recommends using Sqoop 1. Upgrading Sqoop 2 from an earlier release to CDH 5.2.0 and later entails a schema upgrade to the repository database. For details, see [Upgrading Sqoop 2](#).

Apache Whirr Incompatible Changes and Limitations

The Apache Software Foundation has voted to terminate the Whirr project. Whirr is deprecated in CDH 5 and will be removed altogether in a future release.

Apache ZooKeeper Incompatible Changes and Limitations

There are no known incompatible changes between CDH 4 and CDH 5.

Known Issues in CDH 5

Operating System Known Issues

Known issues and workarounds related to operating systems are listed below.

Kerberos credentials cache issue on keyring-enabled configurations

Some operating system configurations use keyring as the default storage for the Kerberos credential cache. This configuration may cause authentication failures for various components, indicated by the error:

```
GSS initiate failed [Caused by GSSEException: No valid credentials provided (Mechanism level: Failed to find any Kerberos tgt)]
```

Affected OS/version: 5.x

Cloudera Bug: CDH-70529

Workaround: Use a file-based Kerberos credentials cache configuration in `krb5.conf`. For details, refer to [SUSE Linux Enterprise Server 12 \(SLES12\) and Kerberos Credential Cache Issue](#) on page 1947.

YARN failure due to Debian 8 kernel limitation

The Linux kernel for Debian 8 (Jessie) does not support the `cgroup` setting for `cpu.cfs_period_us` and causes failures such as that shown in the log message below:

```
...DATE TIME FATAL NodeManager
Error starting NodeManager
org.apache.hadoop.yarn.exceptions.YarnRuntimeException: Failed to initialize container
executor
at
org.apache.hadoop.yarn.server.nodemanager.NodeManager.serviceInit(NodeManager.java:269)
at org.apache.hadoop.service.AbstractService.init(AbstractService.java:163)
...
```

Bug: [Debian 805498](#)

Cloudera Bug: CDH-57554

Linux kernel security patch and CDH services crashes

After applying a recent Linux kernel security patch for [CVE-2017-1000364](#), CDH services that use the JSVC set of libraries crash with a Java Virtual Machine (JVM) error such as:

```
A fatal error has been detected by the Java Runtime Environment:
SIGBUS (0x7) at pc=0x00007fe91ef6cebc, pid=30321, tid=0x00007fe930c67700
```

Cloudera services for HDFS and Impala cannot start after applying the patch.



Important: If you have not upgraded your Linux kernel using the distribution's patch for CVE-2017-1000364, do **not** apply the patch.

Commonly used Linux distributions are shown in the table below. However, the issue affects any CDH release that runs on RHEL, CentOS, Oracle Linux, SUSE Linux, or Ubuntu and that has had the Linux kernel security patch for CVE-2017-1000364 applied.

If you have already applied the patch for your OS according to the advisories for CVE-2017-1000364, apply the kernel update that contains the fix for your operating system (some of which are listed in the table). If you cannot apply the kernel update, you can workaround the issue by increasing the Java thread stack size as detailed in the steps below.

Distribution	Advisories for CVE-2017-1000364	Advisory updates
Oracle Linux 6	ELSA-2017-1486	Oracle has fixed this problem in ELSA-2017-1723 .
Oracle Linux 7	ELSA-2017-1484	Oracle has also added the fix for Oracle Linux 7 in ELBA-2017-1674 .
RHEL 6	RHSA-2017-1486	RedHat has fixed this problem for RHEL 6, marked this as outdated and superseded by RHSA-2017-1723 .
RHEL 7	RHSA-2017-1484	RedHat has fixed this problem for RHEL 7 and has marked this patch as outdated and superseded by RHBA-2017-1674 .
SUSE	CVE-2017-1000364	SUSE has also fixed this problem and the patch names are included in this same advisory.
Ubuntu 14.04	USN-3335-1	Fixed in USN-3338-2 .

Workaround

If you cannot apply the kernel update, you can set the Java thread stack size to `-Xss1280k` for the affected services using the appropriate Java configuration option or the environment advanced configuration snippet, as detailed below.

For role instances that have specific Java configuration options properties:

1. Log in to Cloudera Manager Admin Console.
2. Select **Clusters > Impala**, and then click the **Configuration** tab.
3. Type `java` in the search field to display Java related configuration parameters. The **Java Configuration Options for Catalog Server** property field displays. Type `-Xss1280k` in the entry field, adding to any existing settings.
4. Click **Save Changes**.
5. Navigate to the HDFS service by selecting **Clusters > HDFS**.
6. Click the **Configuration** tab.
7. Click the Scope filter **DataNode**. The **Java Configuration Options for DataNode** field displays among the properties listed. Enter `-Xss1280k` into the field, adding to any existing properties.
8. Click **Save Changes**.
9. Select the Scope filter **NFS Gateway**. The **Java Configuration Options for NFS Gateway** field displays among the properties listed. Enter `-Xss1280k` into the field, adding to any existing properties.
10. Click **Save Changes**.
11. Restart the affected roles (or configure the safety valves in next section and restart when finished with all configurations).

For role instances that do not have specific Java configuration options:

1. Log in to Cloudera Manager Admin Console.
2. Select **Clusters > Impala**, and then click the **Configuration** tab.
3. Click the Scope filter **Impala Daemon** and Category filter **Advanced**.
4. Type `impala daemon environment` in the search field to find the safety valve entry field.
5. In the **Impala Daemon Environment Advanced Configuration Snippet (Safety Valve)**, enter:

```
JAVA_TOOL_OPTIONS=-Xss1280K
```

6. Click **Save Changes**.
7. Click the Scope filter **Impala StateStore** and Category filter **Advanced**.
8. In the **Impala StateStore Environment Advanced Configuration Snippet (Safety Valve)**, enter:

```
JAVA_TOOL_OPTIONS=-Xss1280K
```

9. Click **Save Changes**.

10 Restart the affected roles.

The table below summarizes the parameters that can be set for the affected services:

Service	Settable Java Configuration Option
HDFS DataNode	Java Configuration Options for DataNode
HDFS NFS Gateway	Java Configuration Options for NFS Gateway
Impala Catalog Server	Java Configuration Options for Catalog Server
Impala Daemon	Impala Daemon Environment Advanced Configuration Snippet (Safety Valve)
	JAVA_TOOL_OPTIONS=-Xss1280K
Impala StateStore	Impala StateStore Environment Advanced Configuration Snippet (Safety Valve)
	JAVA_TOOL_OPTIONS=-Xss1280K

Cloudera Bug: CDH-55771


Some OS versions reject weak ciphers for TLS

For some OS configurations, the `curl` command cannot connect to certain CDH services that run on Apache Tomcat when the cluster has been configured for TLS/SSL. Specifically, HttpFS, KMS, Oozie, and Solr services reject connection attempts because the default cipher configuration uses weak temporary server keys (based on Diffie-Hellman key exchange protocol).

Affected OS/version: Linux distributions that use network security services (nss) library version 3.28.4-3.el6_9, including Oracle Linux 6.x, RHEL 6.x, CentOS 6.x. These conflict with Tomcat 6.0.48, which is included with CDH 5.11. The issue has been resolved in CDH 5.11.1.

Workaround: Override the default cipher configuration using all the ciphers from the list below. Only commas separate each cipher, and there are no line endings or other characters in this list.

- For clusters not managed by Cloudera Manager, enter this list in the appropriate configuration file for the service (see the "Not managed" column in the table below). Restart the service.



- For clusters managed by Cloudera Manager, [follow the steps below](#).

Service	Cluster managed by Cloudera Manager		Not managed
	Configuration property	Environment variable	Configuration file
HttpFS	HttpFS Environment Advanced Configuration Snippet (Safety Valve)	HTTPFS_SSL_CIPHERS	/etc/hadoop-httpfs/conf/httpfs-env.sh
KMS	Key Management Server Environment Advanced Configuration Snippet (Safety Valve)	KMS_SSL_CIPHERS	/etc/hadoop-kms/conf/kms-env.sh
Oozie	Oozie Server Environment Advanced Configuration Snippet	OOZIE_HTTPS_CIPHERS	/etc/oozie/conf/oozie-env.sh
Solr	Solr Server Environment Advanced Configuration Snippet	SOLR_CIPHERS_CONFIG	/etc/default/solr

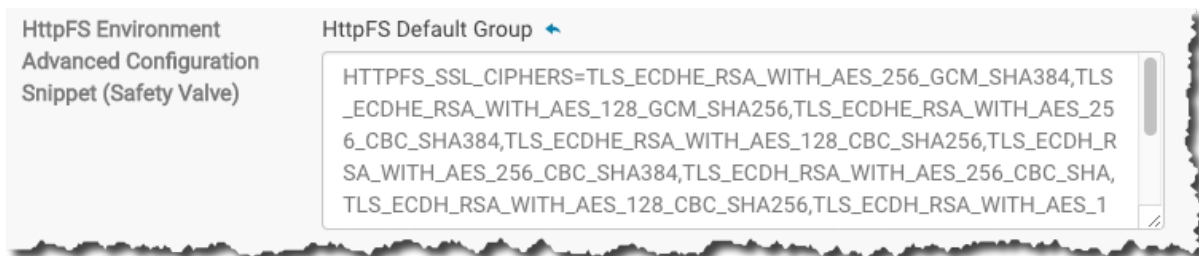
To override the ciphers on Cloudera Manager clusters:

1. Log into Cloudera Manager Admin Console.
2. Search for the advanced configuration snippet name ("Configuration property") for each respective service role.

3. Enter the appropriate environment variable in the entry field.
4. Copy the list from here:

```
HTTPEFS_SSL_CIPHERS=TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,TLS_ECDH_RSA_WITH_AES_256_CBC_SHA,TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256,TLS_ECDH_RSA_WITH_AES_1
```

5. Paste as the setting for the environment variable. For example, overriding HttpFS is as follows:



6. Restart the service.

Cloudera Bugs: CDH-52109, CDH-49236

SUSE Linux Enterprise Server 12 (SLES12) and Kerberos Credential Cache Issue

Cloudera has become aware of a conflict in the handling of the Kerberos credential cache in SLES12 that can result in the following error:

```
Credential cache directory /run/user/uid does not exist
```

The issue is caused by a change in SLES12 `systemd` and the location of Kerberos credentials cache for user accounts. This issue may emerge when password-less accounts, such as `impala` or `flume`, for example, are executed using `sudo -u user` or `su - user` and the system attempts to retrieve the credential from the cache.

Workaround:

To work around this issue, add the following line to your system's Kerberos configuration file (`/etc/krb5.conf`), in the `[libdefaults]` section:

```
[libdefaults]
default_ccache_name = FILE:/tmp/krb5cc_%{uid}
...
```

For clusters managed by Cloudera Manager Server:

1. Log in to Cloudera Manager Admin Console.
2. Click the **Administration** tab and select **Settings** from the drop-down menu.
3. Select **Kerberos** for the **Category** filter.
4. Type the word "safety" in the Search field to find the Advanced Configuration Snippet (Safety Valve) settings for Kerberos.
5. Enter the name=value pair for the property as shown here:



6. Click **Save Changes**.

Cloudera Bugs: OPSAPS-34167, OPSAPS-41158

Random JVM Hangs on RHEL 6.6 (Kernels 3.14-3.17)

User Java processes—HBase RegionServers, HDFS DataNodes—can deadlock and hang for no apparent reason.

Impact: A Java application or daemon hangs. A `jstack`, or attaching a debugger, can unblock the process.

Workaround:

Upgrade the operating system to one with kernel version 3.18 or later; for example, RHEL 6.6z and later. For information on checking your RHEL kernel version, see <http://www.cyberciti.biz/faq/centos-redhat-rhel-6-kernel-version/>. In RHEL, the issue has been fixed recently in kernel-2.6.32-504.16.2.el6 update (April 21, <https://rhn.redhat.com/errata/RHSA-2015-0864.html>). The following example shows how to check for the fix:

```
% rpm -qp --changelog kernel-2.6.32-504.16.2.el6.x86_64.rpm | grep 'Ensure  
get_futex_key_refs() always implies a barrier'
```

Other distributions running kernel versions 3.14-3.17 may be affected.

For more information, see <https://groups.google.com/forum/#!topic/mechanical-sympathy/QbmpZxp6C64>

Cloudera Bug: TSB-63

Leap-Second Events



Note: The [next leap-second event](#) is unknown at this time. The last leap-second event occurred on December 16, 2016 at 23:59:60 UTC.

Impact: After a leap-second event, Java applications (including CDH services) using older Java and Linux kernel versions, may consume almost 100% CPU. See <https://access.redhat.com/articles/15145>.

Leap-second events are tied to the time synchronization methods of the Linux kernel, the Linux distribution and version, and the Java version used by applications running on affected kernels.

Although Java is increasingly agnostic to system clock progression (and less susceptible to a kernel's mishandling of a leap-second event), using JDK 7 or 8 should prevent issues at the CDH level (for CDH components that use the Java Virtual Machine).

Immediate action required:

(1) Ensure that the kernel is up to date.

- **RHEL5/6/7, CentOS 5/6/7** - 2.6.32-298 or higher
- **Ubuntu Trusty 14.04 LTS** - No action required
- **Ubuntu Precise 12.04 LTS** - 3.2.0-29.46 or higher
- **Debian Jessie (8.x)** - No action required
- **Debian Wheezy (7.x)** - 3.2.29-1 or higher
- **Oracle Enterprise Linux (OEL)** - Kernels built in 2013 or later
- **SLES11 (SP2, SP3 and SP4)**
 - **SP2** - 3.0.38-0.5 or higher
 - **SP3 and SP4** - No action required
- **SLES12** - No action required.

(2) Ensure that your Java JDKs are current (especially if the kernel is not up to date and cannot be upgraded).

- **Java 7** - At least `jdk-1.7.0_65`
- **Java 8** - No action required.

(3) Ensure that your systems use either NTP or PTP synchronization.

For systems not using time synchronization, update both the OS tzdata and Java tzdata packages to the tzdata-2016g version, at a minimum. For OS tzdata package updates, contact OS support or check updated OS repositories. For Java tzdata package updates, see Oracle's [Timezone Updater Tool](#).

Cloudera Bugs: CDH-44788, TSB-189

Performance Known Issues



Important: For best practices, and solutions to known performance problems, see [Optimizing Performance in CDH](#).

Install and Upgrade Known Issues

[Upgrade to CDH 5.13 or higher Requires Pre-installation of Spark 2.1 or Spark 2.2](#)

If your cluster has Spark 2.0 or Spark 2.1 installed and you want to upgrade to CDH 5.13 or higher, you must first upgrade to Spark 2.1 release 2 or later before upgrading CDH. To install these versions of Spark, do the following before running the CDH Upgrade Wizard:

1. Install the Custom Service Descriptor (CSD) file. See

- [Installing Spark 2.1](#)
- [Installing Spark 2.2](#)



Note:

Spark 2.2 requires that JDK 1.8 be deployed throughout the cluster. JDK 1.7 is not supported for Spark 2.2.

See [Step 2: Install Java Development Kit](#).

2. Download, distribute, and activate the Parcel for the version of Spark that you are installing:

- **Spark 2.1 release 2:** The parcel name includes "cloudera2" in its name.
- **Spark 2.2 release 1:** The parcel name includes "cloudera1" in its name.

See [Managing Parcels](#).

Affected versions: CDH 5.13.0 and higher

Cloudera bug: CDH-56775

Flume Kafka client incompatible changes in CDH 5.8

Due to the change of offset storage from ZooKeeper to Kafka in the CDH 5.8 Flume Kafka client, data might not be consumed by the Flume agents, or might be duplicated (if `kafka.auto.offset.reset=smallest`) during an upgrade to CDH 5.8.

Workaround: See [Upgrading to CDH 5.8 When Using the Flume Kafka Client](#)

Cloudera Bug: TSB-173

Upgrades to CDH 5.4.1 from Releases Earlier than 5.4.0 May Fail

Problem: Because of a change in the implementation of the NameNode metadata upgrade mechanism, upgrading to CDH 5.4.1 from a version lower than 5.4.0 can take an inordinately long time. In a cluster with NameNode high availability (HA) configured and a large number of edit logs, the upgrade can fail, with errors indicating a timeout in the pre-upgrade step on JournalNodes.

What to do:

To avoid the problem: Do not upgrade to CDH 5.4.1; upgrade to CDH 5.4.2 instead.

If you experience the problem: If you have already started an upgrade and seen it fail, contact Cloudera Support. This problem involves no risk of data loss, and manual recovery is possible.

If you have already completed an upgrade to CDH 5.4.1, or are installing a new cluster: In this case you are not affected and can continue to run CDH 5.4.1.

Potential job failures during YARN rolling upgrades to CDH 5.3.4

Problem: A MapReduce security fix introduced a compatibility issue that results in job failures during YARN rolling upgrades from CDH 5.3.3 to CDH 5.3.4.

Cloudera Bug: CDH-28680

Release affected: CDH 5.3.4

Release containing the fix: CDH 5.3.5

Workarounds: You can use any one of the following workarounds for this issue:

- Upgrade to CDH 5.3.5.
- Restart any jobs that might have failed during the upgrade.
- Explicitly set the version of MapReduce to be used so it is picked on a per-job basis.
 1. Update the YARN property, **MR Application Classpath** (`mapreduce.application.classpath`), either in Cloudera Manager or in the `mapred-site.xml` file. Remove all existing values and add a new entry: `<parcel-path>/lib/hadoop-mapreduce/*`, where `<parcel-path>` is the absolute path to the parcel installation. For example, the default installation path for the CDH 5.3.3 parcel would be: `/opt/cloudera/parcels/CDH-5.3.3-1.cdh5.3.3.p0.5/lib/hadoop-mapreduce/*`.
 2. Wait until jobs submitted with the above client configuration change have run to completion.
 3. Upgrade to CDH 5.3.4.
 4. Update the **MR Application Classpath** (`mapreduce.application.classpath`) property to point to the new CDH 5.3.4 parcel.

Do *not* delete the old parcel until after all jobs submitted prior to the upgrade have finished running.

NameNode Incorrectly Reports Missing Blocks During Rolling Upgrade

Problem: During a rolling upgrade to any of the releases listed below, the NameNode may report missing blocks after rolling back multiple DataNodes. This is caused by a race condition with block reporting between the DataNode and the NameNode. No permanent data loss occurs, but data can be unavailable for up to six hours before the problem corrects itself.

Releases affected: CDH 5.0.6, 5.1.5, 5.2.5, 5.3.3, 5.4.1, 5.4.2

What to do:

To avoid the problem: Cloudera advises skipping the affected releases and installing a release containing the fix. For example, do not upgrade to CDH 5.4.2; upgrade to CDH 5.4.3 instead.

The releases containing the fix are: CDH 5.3.4, 5.4.3

If you have already completed an upgrade to an affected release, or are installing a new cluster: You can continue to run the release, or upgrade to a release that is not affected.

Upgrading to CDH 5.4 or later requires an HDFS upgrade

Upgrading to CDH 5.4.0 or later from an earlier CDH 5 release requires an HDFS upgrade, and upgrading from a release earlier than CDH 5.2.0 requires additional steps. See [Upgrading from an Earlier CDH 5 Release to the Latest Release](#) for further information. See also [What's New In CDH 5.4.x](#) on page 1838.

JDK Does Not Have Up-to-date tzdata

The Europe/Moscow time zone has changed. This time zone change is not incorporated into JDK lower than 8u31 and 7u75. Cloudera ships JDK7u67, which does not reflect these `tzdata` updates. If you run UDF related to time zone conversion, you receive incorrect results.

Workaround: Upgrade to JDK version 8u31 or 7u75.

Extra step needed on Ubuntu Trusty if you add the Cloudera repository

If you install or upgrade CDH on Ubuntu Trusty using the command line, and add the Cloudera repository yourself (rather than using the "1-click Install" method) you need to perform an additional step to ensure that you get the CDH version of ZooKeeper, rather than the version that is bundled with Trusty.

Cloudera Bug: CDH-23569

Upgrading `hadoop-kms` from 5.2.x and 5.3.x releases fails on SLES

Upgrading `hadoop-kms` fails on SLES when you try to upgrade an existing version from 5.2.x releases earlier than 5.2.4, and from 5.3.x releases earlier than 5.3.2.

Cloudera Bug: CDH-25053

Must build native libraries when installing from tarballs

When installing Hadoop from Cloudera tarballs, you must build your own native libraries. The tarballs do not include libraries that are built for the different distributions and architectures.

Installing from tarballs is deprecated and will be removed from CDH in a future release.

Cloudera Bugs: CDH-7304, CDH-5623

Apache Avro Known Issues

Avro Null Defaults Do Not Work in Partitioned Hive Tables

When querying Hive tables created with Avro recursive schemas, Hive silently drops recursive fields. Any attempt to serialize data which exploits recursive schemas results in dropped fields and throws an exception.

Affected Versions: For all versions of CDH 5, Hive does not support Avro recursive schemas in Hive tables.

Cloudera Bug: CDH-72545

Bug: [HIVE-10187](#)

Workaround: To workaround this issue, modify the Avro schema by removing the recursive section.

Addressed in: Avro recursive schemas for Hive tables are supported in CDH 6 and later.

Apache Flume Known Issues

Flume Kafka client incompatible changes in CDH 5.8

Due to the change of offset storage from ZooKeeper to Kafka in the CDH 5.8 Flume Kafka client, data might not be consumed by the Flume agents, or might be duplicated (if `kafka.auto.offset.reset=smallest`) during an upgrade to CDH 5.8.

Cloudera Bug: TSB-173

Workaround: See [Upgrading to CDH 5.8 When Using the Flume Kafka Client](#)

Hive sink support

Flume does not provide a native sink that stores the data that can be directly consumed by Hive.

Bug: [FLUME-1008](#)

Cloudera Bug: CDH-3912

Workaround: None

Fast Replay does not work with encrypted File Channel

If an encrypted file channel is set to use fast replay, the replay will fail and the channel will fail to start.

Bug: [FLUME-1885](#)

Cloudera Bug: CDH-8017, CDH-10115.

Workaround: Disable fast replay for the encrypted channel by setting `use-fast-replay` to false.

Apache Hadoop Known Issues

This page includes known issues and related topics, including:

Deprecated Properties

Several Hadoop and HDFS properties have been deprecated as of Hadoop 2.0.0 (Hadoop 0.23.1, CDH 4 Beta) and later. For details, see [Deprecated Properties](#).

Hadoop Common

KMS Load Balancing Provider Fails to invalidate Cache on Key Delete

The KMS Load balancing Provider has not been correctly invalidating the cache on key delete operations. The failure to invalidate the cache on key delete operations can result in the possibility that data can be leaked from the framework for a short period of time based on the value of the `hadoop.kms.current.key.cache.timeout.ms` property. Its default value is 30,000ms. When the KMS is deployed in an HA pattern the `KMSLoadBalancingProvider` class will only send the delete operation to one KMS role instance in a round-robin fashion. The code lacks a call to invalidate the cache across all instances and can leave key information including the metadata and key stored (the deleted key) in the cache on one or more KMS instances up to the key cache timeout.

Apache Issue:

- [HADOOP-17208](#)
- [HADOOP-17304](#)

Products affected:

- CDH
- HDP
- CDP

Releases affected:

- CDH 5.x
- CDH 6.x
- CDP 7.0.x
- CDP 7.1.4 and earlier
- HDP 2.6 and later

Users affected: Customers with Data-at-rest encryption enabled that have more than 1 kms role instance and the services Key Cache enabled.

Impact: Key Meta-data and Key material may remain active within the service cache.

Severity: Medium

Action required:

- CDH customers: Upgrade to CDP 7.1.5 or request a patch
- HDP customers: Request a patch

Knowledge article: For the latest update on this issue see the corresponding Knowledge article: [TSB 2020-434: KMS Load Balancing Provider Fails to invalidate Cache on Key Delete](#)

Zip Slip Vulnerability

“Zip Slip” is a widespread arbitrary file overwrite critical vulnerability, which typically results in remote command execution. It was discovered and responsibly disclosed by the Snyk Security team ahead of a public disclosure on June 5, 2018, and affects thousands of projects.

Cloudera has analyzed our use of zip-related software, and has determined that only Apache Hadoop is vulnerable to this class of vulnerability in CDH 5. This has been fixed in upstream Hadoop as CVE-2018-8009.

Products affected: Hadoop

Releases affected:

- CDH 5.12.x and all prior releases
- CDH 5.13.0, 5.13.1, 5.13.2, 5.13.3
- CDH 5.14.0, 5.14.2, 5.14.3
- CDH 5.15.0

Users affected: All

Date of detection: April 19, 2018

Detected by: Snyk

Severity: High

Impact: Zip Slip is a form of directory traversal that can be exploited by extracting files from an archive. The premise of the directory traversal vulnerability is that an attacker can gain access to parts of the file system outside of the target folder in which they should reside. The attacker can then overwrite executable files and either invoke them remotely or wait for the system or user to call them, thus achieving remote command execution on the victim’s machine. The vulnerability can also cause damage by overwriting configuration files or other sensitive resources, and can be exploited on both client (user) machines and servers.

CVE: CVE-2018-8009

Immediate action required: Upgrade to a version that contains the fix.

Addressed in release/refresh/patch: CDH 5.14.4, CDH 5.15.1

For the latest update on this issue, see the corresponding Knowledge article:

TSB: 2018-307: Zip Slip Vulnerability

[Apache Hadoop MapReduce Job History Server \(JHS\) vulnerability CVE-2017-15713](#)

A vulnerability in Hadoop’s Job History Server allows a cluster user to expose private files owned by the user running the MapReduce Job History Server (JHS) process. See <http://seclists.org/oss-sec/2018/q1/79> for reference.

Products affected: Apache Hadoop MapReduce

Releases affected: All releases prior to CDH 5.12.0. CDH 5.12.0, CDH 5.12.1, CDH 5.12.2, CDH 5.13.0, CDH 5.13.1, CDH 5.14.0

Users affected: Users running the MapReduce Job History Server (JHS) daemon

Date/time of detection: November 8, 2017

Detected by: Man Yue Mo of lgtm.com

Severity (Low/Medium/High): High

Impact: The vulnerability allows a cluster user to expose private files owned by the user running the MapReduce Job History Server (JHS) process. The malicious user can construct a configuration file containing XML directives that reference sensitive files on the MapReduce Job History Server (JHS) host.

CVE: CVE-2017-15713

Immediate action required: Upgrade to a release where the issue is fixed.

Addressed in release/refresh/patch: CDH 5.13.2, 5.14.2

[Hadoop LdapGroupsMapping does not support LDAPS for self-signed LDAP server](#)

Hadoop LdapGroupsMapping does not work with LDAP over SSL (LDAPS) if the LDAP server certificate is self-signed. This use case is currently not supported even if Hadoop User Group Mapping LDAP TLS/SSL Enabled, Hadoop User Group Mapping LDAP TLS/SSL Truststore, and Hadoop User Group Mapping LDAP TLS/SSL Truststore Password are filled properly.

Bug: [HADOOP-12862](#)

Affected Versions: All CDH 5 versions.

Workaround: None.

HDFS

[CVE-2018-1296 Permissive Apache Hadoop HDFS listXAttr Authorization Exposes Extended Attribute Key/Value Pairs](#)

AHDFS exposes extended attribute key/value pairs during listXAttrs, verifying only path-level search access to the directory rather than path-level read permission to the referent.

Products affected: Apache HDFS

Releases affected:

- CDH 5.4.0 - 5.15.1, 5.16.0
- CDH 6.0.0, 6.0.1, 6.1.0

Users affected: Users who store sensitive data in extended attributes, such as users of HDFS encryption.

Date/time of detection: December 12, 2017

Detected by: Rushabh Shah, Yahoo! Inc., Hadoop committer

Severity (Low/Medium/High): Medium

Impact: HDFS exposes extended attribute key/value pairs during listXAttrs, verifying only path-level search access to the directory rather than path-level read permission to the referent. This affects features that store sensitive data in extended attributes.

CVE: CVE-2018-1296

Immediate action required:

- Upgrade: Update to a version of CDH containing the fix.
- Workaround: If a file contains sensitive data in extended attributes, users and admins need to change the permission to prevent others from listing the directory that contains the file.

Addressed in release/refresh/patch:

- CDH 5.15.2, 5.16.1
- CDH 6.1.1, 6.2.0

Clusters running CDH 5.16.1, 6.1.0, or 6.1.1 can lose some HDFS file permissions any time the NameNode is restarted

When a cluster is upgraded to 5.16.1, 6.1.0, or 6.1.1 roles with SELECT and/or INSERT privileges on an Impala database or table will have the REFRESH privilege added as part of the upgrade process. HDFS ACLs for roles with the REFRESH privilege get set with empty permissions whenever the NameNode is restarted. This can cause any jobs or queries run by users within affected roles to fail because they will no longer be able to access affected Impala database or tables.

Products Affected: HDFS and components that access files in HDFS

Affected Versions: CDH 5.16.1, 6.1.0, 6.1.1

Users Affected: Clusters with Impala and HDFS ACLs managed by Sentry upgrading from any release to CDH 5.16.1, 6.1.0, and 6.1.1.

Severity (Low/Medium/High): High

Root Cause and Impact: The new privilege REFRESH was introduced in CDH 5.16 and 6.1 and applies to Impala databases and tables. When a cluster is upgraded to 5.16.1, 6.1.0, or 6.1.1, roles with SELECT or INSERT privileges on an Impala database or table will have the REFRESH privilege added during the upgrade.

HDFS ACLs for roles with the REFRESH privilege get set with empty permissions whenever the NameNode is restarted. The NameNode is restarted during the upgrade.

For example if a group `appdev` is in role `appdev_role` and has SELECT access to the Impala table "project" the HDFS ACLs prior to the upgrade would look similar to:

```
group: appdev
group: :r--
```

After the upgrade the HDFS ACLs will be set with no permissions and will look like this:

```
group: appdev
group: :---
```

Any jobs or queries run by users within affected roles will fail because they will no longer be able to access affected Impala database or tables. This impacts any SQL client accessing the affected databases and tables. For example, if a Hive client is used to access a table created in Impala it will also fail. Jobs accessing the files directly through HDFS, e.g. via Spark, will also be impacted.

The HDFS ACLs will get reset whenever the NameNode is restarted.

Immediate action required: If possible, do not upgrade to releases CDH 5.16.1, 6.1.0, or 6.1.1 if Impala is used and Sentry manages HDFS ACLs within your environment. Subsequent CDH releases will resolve the problem with a product fix under [SENTRY-2490](#).

If an upgrade is being considered, reach out to your account team to discuss other possibilities, and to receive additional insight into future product release schedules.

If an upgrade must be executed, contact Cloudera Support indicating the upgrade plan and why an upgrade is being executed. Options are available to assist with the upgrade if necessary.

Addressed in release/refresh/patch: Patches for 5.16.1, 6.1.0 and 6.1.1 are available for major supported operating systems. Customers are encouraged to contact Cloudera Support for a patch. The patch should be applied immediately after upgrade to any of the affected versions.

The fix for this TSB will be included in 6.1.2, 6.2.0, 5.16.2, and 5.17.0.

Potential data corruption due to race conditions between concurrent block read and write

Under rare conditions when an HDFS file is open for write, an application reading the same HDFS blocks might read up-to-date block data of the partially written file, while reading a stale checksum that corresponds to the block data before the latest write. The block is incorrectly declared corrupt as a result. Normally the HDFS NameNode schedules additional replica for the same block from other replicas if a replica is corrupted, but if the frequency of concurrent write and read is high enough, there is a small probability that all replicas of a block can be declared corrupt, and the file becomes corrupt and unrecoverable as well.

DataNode might print error log such as the following:

```
2017-10-18 11:23:46,627 ERROR org.apache.hadoop.hdfs.server.datanode.DataNode:
ip-168-61-2-30:50010:DataXceiver error processing WRITE_BLOCK operation src:
/168.61.2.32:48163 dst: /168.61.12.31:50010
java.io.IOException: Terminating due to a checksum error.java.io.IOException: Unexpected
checksum mismatch while writing
BP-1666924250-168.61.12.36-1494235758065:blk_1084584428_5057054 from /168.61.12.32:48163
at
org.apache.hadoop.hdfs.server.datanode.BlockReceiver.receivePacket(BlockReceiver.java:604)
at
org.apache.hadoop.hdfs.server.datanode.BlockReceiver.receiveBlock(BlockReceiver.java:894)
at org.apache.hadoop.hdfs.server.datanode.DataXceiver.writeBlock(DataXceiver.java:794)
at org.apache.hadoop.hdfs.protocol.datatransfer.Receiver.opWriteBlock(Receiver.java:169)
at org.apache.hadoop.hdfs.protocol.datatransfer.Receiver.processOp(Receiver.java:106)
at org.apache.hadoop.hdfs.server.datanode.DataXceiver.run(DataXceiver.java:246)
at java.lang.Thread.run(Thread.java:745)
```

The bug is fixed by HDFS-11056, HDFS-11160 and HDFS-11229.

Products Affected: HDFS

Affected Versions:

- All CDH 5.4 releases and lower
- CDH 5.5.0, 5.5.1, 5.5.2, 5.5.4, 5.5.5
- CDH 5.6.0, 5.6.1
- CDH 5.7.0, 5.7.1, 5.7.2, 5.7.3, 5.7.4, 5.7.5
- CDH 5.8.0, 5.8.2, 5.8.3
- CDH 5.9.0, 5.9.1

Users Affected: Workloads that require reading a file while it's being concurrently written to HDFS.

Severity (Low/Medium/High): Low

Impact: If the workload requires reading and writing the same file concurrently, there is a small probability that all replicas of a block can be declared corrupt, and the file becomes corrupt as well.

Immediate action required: Customers are advised to upgrade to a CDH version containing the fix if the workloads are susceptible to this bug.

Fixed in Versions:

- CDH 5.5.6 and higher
- CDH 5.7.6 and higher
- CDH 5.8.4 and higher
- CDH 5.9.2 and higher
- CDH 5.10.0 and higher

[Cannot re-encrypt an encryption zone if a previous re-encryption on it was canceled](#)

When canceling a re-encryption on an encryption zone, the status of the re-encryption may continue to show "Processing". When this occurs, future re-encrypt commands for this encryption zone will fail inside the NameNode, and the re-encryption will never complete.

Cloudera Bug: CDH-59073

Affected Versions: CDH 5.13.0

Fixed in Versions: CDH 5.13.1 and higher

Workaround: To halt, or remove the "Processing" status for the encryption zone, re-issue the cancel re-encryption command on the encryption zone. If a new re-encryption command is required for this encryption zone, restart the NameNode before issuing the command.

Potential Block Corruption and Data Loss During Pipeline Recovery

A bug in the HDFS block pipeline recovery code can cause blocks to be unrecoverable due to miscalculation of the block checksum. On a busy cluster where data is written and flushed frequently, when a write pipeline recovery occurs, a node newly added to the write pipeline may calculate the checksum incorrectly. This miscalculation is very rare, but when it does occur, the replica becomes corrupted and data can be lost if all replicas are simultaneously affected.

Bug: [HDFS-4660](#), [HDFS-9220](#)

Detecting this known issue requires correlating multiple log messages. Below is an example DataNode log error message captured at the time of block creation:

```
java.io.IOException: Terminating due to a checksum error.java.io.IOException:
Unexpected checksum mismatch while writing
BP-1800173197-10.x.y.z-1444425156296:blk_1170125248_96458336 from /10.x.y.z
```



Important: Absence of error messages in the DataNode logs does not necessarily mean that the issue did not occur.

This issue affects these versions of CDH:

- CDH 5.0.0, 5.0.1, 5.0.2, 5.0.3, 5.0.4, 5.0.5, 5.0.6
- CDH 5.1.0, 5.1.2, 5.1.3, 5.1.4, 5.1.5
- CDH 5.2.0, 5.2.1, 5.2.3, 5.2.4, 5.2.5, 5.2.6
- CDH 5.3.0, 5.3.1, 5.3.2, 5.3.3, 5.3.4, 5.3.5, 5.3.7, 5.3.8, 5.3.9, 5.3.10
- CDH 5.4.0, 5.4.1, 5.4.2, 5.4.3, 5.4.4, 5.4.5, 5.4.7, 5.4.8, 5.4.9, 5.4.10
- CDH 5.5.0, 5.5.1

Workaround: None. Upgrade to a CDH version that includes the fix: CDH 5.4.11, CDH 5.5.2, CDH 5.6.0 and higher.



Important: Upgrading does not fix existing corrupted block replicas. The HDFS block scanner runs every three weeks and captures "Checksum failed" WARN messages in the DataNode log. Look for these to identify and repair corrupted block replicas.

Users affected: All users running the affected CDH versions and using the HDFS file system.

Severity (Low/Medium/High): High

Impact: Potential loss of block data.

Immediate action required: Upgrade to a CDH version that includes the fix, specifically:

- CDH 5.4.11, CDH 5.5.2, CDH 5.6.0 and higher

DiskBalancer Occasionally Emits False Error Messages

Diskbalancer occasionally emits false error messages. For example:

```
2016-08-03 11:01:41,788 ERROR org.apache.hadoop.hdfs.server.datanode.DiskBalancer:
Disk Balancer is not enabled.
```

You can safely ignore this error message if you are not using DiskBalancer.

Affected Versions: CDH 5.8.1 and below.

Fixed in Versions: CDH 5.8.2 and higher.

Bug: [HDFS-10588](#)

Workaround: Use the following command against all DataNodes to suppress DiskBalancer logs:

```
hadoop daemonlog -setlevel <host:port> org.apache.hadoop.hdfs.server.datanode.DiskBalancer FATAL
```

Another workaround is to suppress the warning by setting the log level of DiskBalancer to FATAL. Add the following to `log4j.properties` (DataNode Logging Advanced Configuration Snippet (Safety Valve)) and restart your DataNodes:

```
log4j.logger.org.apache.hadoop.hdfs.server.datanode.DiskBalancer = FATAL
```

Upgrade Requires an HDFS Upgrade

Upgrading from any release earlier than CDH 5.2.0 to CDH 5.2.0 or later requires an HDFS Upgrade.

Optimizing HDFS Encryption at Rest Requires Newer openssl Library on Some Systems

CDH 5.3 implements the **Advanced Encryption Standard New Instructions** (AES-NI), which provide substantial performance improvements. To get these improvements, you need a recent version of `libcrypto.so` on HDFS and MapReduce client hosts that is, any host from which you originate HDFS or MapReduce requests. Many OS versions have an older version of the library that does not support AES-NI.

See [HDFS Transparent Encryption](#) in the *Encryption* section of the *Cloudera Security* guide for instructions for obtaining the right version.

Other HDFS Encryption Known Issues

Potentially Incorrect Initialization Vector Calculation in HDFS Encryption

A mathematical error in the calculation of the Initialization Vector (IV) for encryption and decryption in HDFS could cause data to appear corrupted when read. The IV is a 16-byte value input to encryption and decryption ciphers. The calculation of the IV implemented in HDFS was found to be subtly different from that used by Java and OpenSSL cryptographic routines. The result is that data could possibly appear to be corrupted when it is read from a file inside an Encryption Zone.

Fortunately, the probability of this occurring is extremely small. For example, the maximum size of a file in HDFS is 64 TB. This enormous file would have a 1-in-4-million chance of hitting this condition. A more typically sized file of 1 GB would have a roughly 1-in-274-billion chance of hitting the condition.

Affected Versions: CDH 5.2.1 and below

Fixed in Versions: CDH 5.3.0 and higher

Cloudera Bug: CDH-23618

Workaround: If you are using the experimental HDFS encryption feature in CDH 5.2, upgrade to CDH 5.3 and verify the integrity of all files inside an Encryption Zone.

DistCp between unencrypted and encrypted locations fails

By default, DistCp compares checksums provided by the filesystem to verify that data was successfully copied to the destination. However, when copying between unencrypted and encrypted locations, the filesystem checksums will not match since the underlying block data is different.

Affected Versions: CDH 5.2.1 and below.

Fixed in Versions: CDH 5.2.2 and higher.

Bug: [HADOOP-11343](#)

Workaround: Specify the `-skipcrccheck` and `-update` distcp flags to avoid verifying checksums.

Cannot move encrypted files to trash

With HDFS encryption enabled, you cannot move encrypted files or directories to the trash directory.

Affected Versions: All CDH 5 versions

Bug: [HDFS-6767](#)

Workaround: To remove encrypted files/directories, use the following command with the `-skipTrash` flag specified to bypass trash.

```
rm -r -skipTrash /testdir
```

HDFS NFS gateway and CDH installation (using packages) limitation

HDFS NFS gateway works as shipped ("out of the box") only on RHEL-compatible systems, but not on SLES, Ubuntu, or Debian. Because of a bug in native versions of `portmap/rpcbind`, the HDFS NFS gateway does not work out of the box on SLES, Ubuntu, or Debian systems when CDH has been installed from the command-line, using packages. It does work on supported versions of RHEL-compatible systems on which `rpcbind-0.2.0-10.el6` or later is installed, and it does work if you use Cloudera Manager to install CDH, or if you start the gateway as root.

Bug: [731542](#) (Red Hat), [823364](#) (SLES), [594880](#) (Debian)

Workarounds and caveats:

- On Red Hat and similar systems, make sure `rpcbind-0.2.0-10.el6` or later is installed.
- On SLES, Debian, and Ubuntu systems, do one of the following:
 - Install CDH using Cloudera Manager; *or*
 - As of CDH 5.1, start the NFS gateway as root; *or*
 - [Start the NFS gateway without using packages](#); *or*
 - You can use the gateway by running `rpcbind` in insecure mode, using the `-i` option, but keep in mind that this allows anyone from a remote host to bind to the portmap.

HDFS does not currently provide ACL support for the NFS gateway

Affected Versions: All CDH 5 versions

Bug: [HDFS-6949](#)

Cloudera Bug: CDH-26921

No error when changing permission to 777 on `.snapshot` directory

Snapshots are read-only; running `chmod 777` on the `.snapshots` directory does not change this, but does not produce an error (though other illegal operations do).

Affected Versions: All CDH 5 versions

Bug: [HDFS-4981](#)

Cloudera Bug: CDH-13062

Workaround: None

Snapshot operations are not supported by ViewFileSystem

Affected Versions: All CDH 5 versions

Cloudera Bug: CDH-12600

Workaround: None

Snapshots do not retain directories' quotas settings

Affected Versions: All CDH 5 versions

Bug: [HDFS-4897](#)

Workaround: None

Permissions for `dfs.namenode.name.dir` incorrectly set.

Hadoop daemons should set permissions for the `dfs.namenode.name.dir` (or `dfs.name.dir`) directories to `drwx-----` (700), but in fact these permissions are set to the file-system default, usually `drwxr-xr-x` (755).

Affected Versions: All CDH 5 versions

Bug: [HDFS-2470](#)

Workaround: Use `chmod` to set permissions to 700.

`hadoop fsck -move` does not work in a cluster with host-based Kerberos

Affected Versions: All CDH 5 versions

Cloudera Bug: CDH-7017

Workaround: Use `hadoop fsck -delete`

HttpFS cannot get delegation token without prior authenticated request

A request to obtain a delegation token cannot initiate an SPNEGO authentication sequence; it must be accompanied by an authentication cookie from a prior SPNEGO authentication sequence.

Affected Versions: CDH 5.1 and below

Fixed in Versions: CDH 5.2 and higher

Bug: [HDFS-3988](#)

Cloudera Bug: CDH-8144

Workaround: Make another WebHDFS request (such as `GETHOMEDIR`) to initiate an SPNEGO authentication sequence and then make the delegation token request.

`DistCp` does not work between a secure cluster and an insecure cluster in some cases

See the upstream bug reports for details.

Affected Versions: All CDH 5 versions

Bug: [HDFS-7037](#), [HADOOP-10016](#), [HADOOP-8828](#)

Cloudera Bug: CDH-14945, CDH-18779

Workaround: None

Port configuration required for `DistCp` to `Hftp` from secure cluster (SPNEGO)

To copy files using `DistCp` to `Hftp` from a secure cluster using SPNEGO, you must configure the `dfs.https.port` property on the client to use the HTTP port (50070 by default).

Affected Versions: All CDH 5 versions

Bug: [HDFS-3983](#)

Cloudera Bug: CDH-8118

Workaround: Configure `dfs.https.port` to use the HTTP port on the client

Non-HA DFS Clients do not attempt reconnects

This problem means that streams cannot survive a NameNode restart or network interruption that lasts longer than the time it takes to write a block.

Affected Versions: All CDH 5 versions

Bug: [HDFS-4389](#)

Cloudera Bug: CDH-10415

DataNodes may become unresponsive to block creation requests

DataNodes may become unresponsive to block creation requests from clients when the directory scanner is running.

Affected Versions: CDH 5.2.1 and below

Fixed in Versions: CDH 5.2.2 and higher

Bug: [HDFS-7489](#)

Workaround: Disable the directory scanner by setting `dfs.datanode.directoryscan.interval` to `-1`.

The active NameNode will not accept an `fsimage` sent from the standby during rolling upgrade

The result is that the NameNodes fail to checkpoint until the upgrade is finalized.



Note: Rolling upgrade is supported only for clusters managed by Cloudera Manager; you cannot do a rolling upgrade in a command-line-only deployment.

Affected Versions: CDH 5.3.7 and below

Fixed in Versions: CDH 5.3.8 and higher

Bug: [HDFS-7185](#)

Workaround: None.

Block report can exceed maximum RPC buffer size on some DataNodes

On a DataNode with a large number of blocks, the block report may exceed the maximum RPC buffer size.

Affected Versions: All CDH 5 versions

Bug: None

Workaround: Increase the value `ipc.maximum.data.length` in `hdfs-site.xml`:

```
<property>
  <name>ipc.maximum.data.length</name>
  <value>268435456</value>
</property>
```

Misapplied user-limits setting possible

The `ulimits` setting in `/etc/security/limits.conf` is applied to the wrong user when security is enabled.

Affected Versions: CDH 5.2.0 and below

Bug: [DAEMON-192](#)

Anticipated Resolution: None

Workaround: To increase the `ulimits` applied to DataNodes, you must change the `ulimit` settings for the root user, not the `hdfs` user.

LAZY_PERSIST storage policy is experimental and not supported

Using this storage policy could potentially lead to data loss.

Affected versions: All CDH 5 versions

Bug: [HDFS-8229](#)

Workaround: None

MapReduce2, YARN

NodeManager fails because of the changed default location of container executor binary

The default location of container-executor binary and .cfg files was changed to `/var/lib/yarn-ce`. It used to be `/opt/cloudera/parcels/<CDH_parcel_version>`. Because of this change, if you did not have the mount options `-noexec` and `-nosuid` set on `/opt`, the NodeManager can fail to start up as these options are set on `/var`.

Affected versions CDH 5.16.1, All CDH 6 versions

Workaround: Either remove the `-noexec` and `-nosuid` mount options on `/var` or change the container-executor binary and .cdf path using the `CMF_YARN_SAFE_CONTAINER_EXECUTOR_DIR` environment variable.

YARN scheduler queue ACLs are not checked when performing MoveApplicationAcrossQueues operations

The YARN `moveApplicationAcrossQueues` operation does not check ACLs on the target queue. This allows a user to move an application to a queue that the user has no access to.

Affected Versions: All CDH 5 versions

Fixed Versions: CDH 6.0.0

Bug: YARN-5554

Cloudera Bug: CDH-43327

Workaround: N/A

Hadoop YARN Privilege Escalation CVE-2016-6811

A vulnerability in Hadoop YARN allows a user who can escalate to the yarn user the ability to possibly run arbitrary commands as the root user.

Products affected: Hadoop YARN

Releases affected:

- CDH 5.12.x and all prior releases
- CDH 5.13.0, 5.13.1, 5.13.2, 5.13.3
- CDH 5.14.0, 5.14.2, 5.14.3
- CDH 5.15.0

Users affected: Users running the Hadoop YARN service.

Detected by: Freddie Rice

Severity: High

Impact: The vulnerability allows a user who has access to a node in the cluster running a YARN NodeManager and who can escalate to the yarn user, the ability to run arbitrary commands as the root user even if the user is not allowed to escalate directly to the root user.

CVE: CVE-2016-6811

Upgrade: Upgrade to a release where the issue is fixed.

Workaround: The vulnerability can be mitigated by restricting access to the nodes where the YARN NodeManagers are deployed, and by removing su access to the yarn user and by making sure no one other than the yarn user is a member of the yarn group. Please consult with your internal system administration team and adhere to your internal security policy when evaluating the feasibility of the above mitigation steps.

Addressed in release/refresh/patch: CDH 5.14.4, 5.15.1

For the latest update on this issue, see the corresponding Knowledge article:

TSB: 2018-309: Hadoop YARN privilege escalation

Missing results in Hive, Spark, Pig, Custom MapReduce jobs, and other Java applications when filtering Parquet data written by Impala

Apache Hive and Apache Spark rely on Apache Parquet's `parquet-mr` Java library to perform filtering of Parquet data stored in row groups. Those row groups contain statistics that make the filtering efficient without having to examine every value within the row group.

Recent versions of the `parquet-mr` library contain a bug described in [PARQUET-1217](#). This bug causes filtering to behave incorrectly if only some of the statistics for a row group are written. Starting in CDH 5.13, Apache Impala populates statistics in this way for Parquet files. As a result, Hive and Spark may incorrectly filter Parquet data that is written by Impala.

In CDH 5.13, Impala started writing Parquet's `null_count` metadata field without writing the `min` and `max` fields. This is valid, but it triggers the [PARQUET-1217](#) bug in the predicate push-down code of the Parquet Java library (`parquet-mr`). If the `null_count` field is set to a non-zero value, `parquet-mr` assumes that `min` and `max` are also set and reads them without checking whether they are actually there. If those fields are not set, `parquet-mr` reads their default value instead.

For integer SQL types, the default value is 0, so `parquet-mr` incorrectly assumes that the `min` and `max` values are both 0. This causes the problem when filtering data. Unless the value 0 itself matches the search condition, all row groups are discarded due to the incorrect `min/max` values, which leads to missing results.

Affected Products: The Parquet Java library (`parquet-mr`) and by extension, all Java applications reading Parquet files, including, but not limited to:

- Hive
- Spark
- Pig
- Custom MapReduce jobs

Affected Versions:

- CDH 5.13.0, 5.13.1, 5.13.2, and 5.14.0
- CDS 2.2 Release 2 Powered by Apache Spark and earlier releases on CDH 5.13.0 and later

Who Is Affected: Anyone writing Parquet files with Impala and reading them back with Hive, Spark, or other Java-based components that use the `parquet-mr` libraries for reading Parquet files.

Severity (Low/Medium/High): High

Impact: Parquet files containing null values for integer fields written by Impala produce missing results in Hive, Spark, and other Java applications when filtering by the integer field.

Immediate Action Required:

- **Upgrade**

You should upgrade to one of the fixed maintenance releases mentioned below.

- **Workaround**

This issue can be avoided at the price of performance by disabling predicate push-down optimizations:

- In Hive, use the following `SET` command:

```
SET hive.optimize.ppd = false;
```

- In Spark, disable the following configuration setting:

```
--conf spark.sql.parquet.filterPushdown=false
```

Addressed in the Following Releases:

- CDH 5.13.3 and higher
- CDH 5.14.2 and higher

- CDH 5.15.0 and higher
- CDS 2.3 Release 2 and higher

For the latest update on this issue, see the corresponding Knowledge Base article:

[TSB:2018-300: Missing results in Hive, Spark, Pig, and other Java applications when filtering Parquet data written by Impala](#)

Apache Hadoop Yarn Fair Scheduler might stop assigning containers when preemption is on

In CDH 5.11.0 the preemption code was updated to improve preemption behavior. Further changes were implemented in CDH 5.11.1 and CDH 5.12.0 to fix a remaining issue as described in [YARN-6432](#) (FairScheduler: Reserve preempted resources for corresponding applications). This fix resulted in two possible side effects:

- A race condition that results in the Fair Scheduler making duplicate reservations. The duplicate reservations are never released and can result in an integer overflow stopping container assignments.
- A possible deadlock in the event processing of the Fair Scheduler. This will stop all updates in the Resource Manager.

Both side effects will ultimately cause the Fair Scheduler to stop processing resource requests.

Without the change from YARN-6432 the resources that are released after being preempted are not reserved for the starved application. This could result in scheduler assigning the preempted container to any application, not just the starved application. If no reservations are made on the node for the starved application preemption will be less effective in solving the resource starvation.

Products affected: YARN

Releases affected: CDH 5.11.1, 5.12.0

Users affected: Users who have YARN configured with the FairScheduler and have turned preemption on.

Severity (Low/Medium/High): Low

Impact: The Resource Manager will accept application but no application will change state or get container assigned and thus progress.

Immediate action required:

- If you have not upgraded to the affected release and preemption in the FairScheduler is in use, avoid upgrading to the affected releases.
- If you have already upgraded to the affected releases, choose from the following options:
 - Upgrade to CDH 5.11.2 or 5.12.1
 - Turn off preemption

Fixed in Versions: CDH 5.11.2 and 5.12.1

Yarn's Continuous Scheduling can cause slowness in Oozie

When Continuous Scheduling is enabled in Yarn, this can cause slowness in Oozie due to long delays in communicating with Yarn. In Cloudera Manager 5.9.0 and higher, `Enable Fair Scheduler Continuous Scheduler` is turned off by default.

Affected Versions: All CDH 5 versions

Bug: None

Cloudera Bug: CDH-60788

Workaround: Turn off `Enable Fair Scheduler Continuous Scheduling` in Cloudera Manager's Yarn Configuration. To keep equivalent benefits of this feature, turn on `Fair Scheduler Assign Multiple Tasks`.

Rolling upgrades to 5.11.0 and 5.11.1 may cause application failures

Affected Versions: CDH versions that can be upgraded to 5.11.0 or 5.11.1

Fixed in Versions: CDH 5.11.2 and higher

Bug: None

Cloudera Bug: CDH-55284, TSB-241

Workaround: Upgrade to 5.11.2 or higher.

[Name resolution issues can result in unresponsive Web UI and REST endpoints](#)

Name resolution issues can cause the Web UI or the RM REST endpoints to consume all ResourceManager request handling threads, leaving the Web UI and REST endpoints unresponsive.

Fixed in Versions: CDH 5.10.0 and higher

Bug: [YARN-4767](#)

Cloudera Bug: CDH-45597

Workaround: Restart the ResourceManager or kill the application that is being accessed or waiting for the ResourceManager to complete the job.

[Loss of connection to the Zookeeper cluster can cause problems with the ResourceManagers](#)

Loss of connection to the Zookeeper cluster can cause the ResourceManagers to be in active-active state for an extended period of time.

Fixed in Versions: CDH 5.10.0 and higher

Bug: [YARN-5677](#), [YARN-5694](#)

Cloudera Bug: CDH-45210

Workaround: None.

[If the YARN user is granted access to all keys in KMS, then files localized from an encryption zone can be world readable](#)

If the YARN user is granted access to all keys in KMS, then files localized from an encryption zone can be world readable.

Fixed in Versions: CDH 5.7.7, 5.8.5, 5.9.2, 5.10.1, 5.11.0 and higher.

Bug: None

Cloudera Bug: CDH-47377

Workaround: Make sure files in an encryption zone do not have world-readable files modes if they are going to be localized.

[Zookeeper outage can cause the ResourceManagers to exit](#)

Fixed in Versions: CDH 5.12.0 and higher

Bug: [YARN-3742](#)

Cloudera Bug: CDH-47439

Workaround: None.

[FairScheduler might not Assign Containers](#)

Under certain circumstances, turning on `Fair Scheduler Assign Multiple Tasks`

(`yarn.scheduler.fair.assignmultiple`) causes the scheduler to stop assigning containers to applications.

Possible symptoms are that running applications show no progress, and new applications do not start, staying in an `Assigned` state, despite the availability of free resources on the cluster.

Affected Versions: CDH 5.5.0, CDH-5.5.1, CDH-5.5.2, CDH-5.5.3, CDH-5.5.4, CDH-5.5.5, CDH-5.5.6, CDH-5.6.0, and CDH-5.6.1

Fixed in Versions: CDH 5.7.0 and higher

Bug: [YARN-4477](#)

Cloudera Bug: CDH-36686

Workaround: Turn off `Fair Scheduler Assign Multiple Tasks` (`yarn.scheduler.fair.assignmultiple`) and restart the ResourceManager.

FairScheduler: AMs can consume all vCores leading to a livelock

When using FAIR policy with the FairScheduler, Application Masters can consume all vCores which may lead to a livelock.

Fixed in Versions: CDH 5.7.3 and higher, except for CDH 5.8.0 and CDH 5.8.1

Bug: [YARN-4866](#)

Cloudera Bug: CDH-37529

Workaround: Use Dominant Resource Fairness (DRF) instead of FAIR; or make sure that the cluster has enough vCores in proportion to the memory.

NodeManager mount point mismatch (YARN)

NodeManager may select a `cgroups` (Linux control groups) mount point that is not accessible to user `yarn`, resulting in failure to start up. The mismatch occurs because YARN uses `cgroups` in mount point `/run/lxcfs/controllers`, while Cloudera Manager typically configures `cgroups` at `/sys/fs/cgroups`. This issue has occurred on Ubuntu 16.04 systems.

Fixed in Versions: CDH 5.11.1 and higher

Bug: [YARN-6433](#)

Cloudera Bug: CDH-52263

Workaround: Look through your YARN logs to identify the failing mount point, and unmount it:

```
$ umount errant_mount_point
```

You can also:

1. `apt-get remove lxcfs`
2. Reboot the node

JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

Affected Versions: All CDH 5 versions.

Workaround: For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated `mapred-site.xml` that references the correct JobHistory Server.

Starting an unmanaged ApplicationMaster may fail

Starting a custom Unmanaged ApplicationMaster may fail due to a race in getting the necessary tokens.

Affected Versions: CDH 5.1.5 and below.

Fixed in Versions: CDH 5.2 and higher.

Bug: [YARN-1577](#)

Cloudera Bug: CDH-17405

Workaround: Try to get the tokens again; the custom unmanaged ApplicationMaster should be able to fetch the necessary tokens and start successfully.

Moving jobs between queues not persistent after restart

CDH 5 adds the capability to move a submitted application to a different scheduler queue. This queue placement is not persisted across ResourceManager restart or failover, which resumes the application in the original queue.

Affected Versions: All CDH 5 versions.

Bug: [YARN-1558](#)

Cloudera Bug: CDH-17408

Workaround: After ResourceManager restart, re-issue previously issued move requests.

Encrypted shuffle may fail (MRv2, Kerberos, TLS)

In MRv2, if the `LinuxContainerExecutor` is used (usually as part of Kerberos security), and `hadoop.ssl.enabled` is set to `true` (see [Configuring Encrypted Shuffle, Encrypted Web UIs, and Encrypted HDFS Transport](#)), then the encrypted shuffle does not work and the submitted job fails.

Affected Versions: All CDH 5 versions.

Bug: [MAPREDUCE-4669](#)

Cloudera Bug: CDH-8036

Workaround: Use encrypted shuffle with Kerberos security without encrypted web UIs, or use encrypted shuffle with encrypted web UIs without Kerberos security.

ResourceManager-to-Application Master HTTPS link fails

In MRv2 (YARN), if `hadoop.ssl.enabled` is set to `true` (use HTTPS for web UIs), then the link from the ResourceManager to the running MapReduce Application Master fails with an HTTP Error 500 because of a PKIX exception.

A job can still be run successfully, and, when it finishes, the link to the job history does work.

Affected Versions: CDH versions before 5.1.0.

Fixed Versions: CDH 5.1.0

Bug: [YARN-113](#)

Cloudera Bug: CDH-8014

Workaround: Do not use encrypted web UIs.

History link in ResourceManager web UI broken for killed Spark applications

When a Spark application is killed, the history link in the ResourceManager web UI does not work.

Workaround: To view the history for a killed Spark application, see the Spark HistoryServer web UI instead.

Affected Versions: All CDH versions

Apache Issue: None

Cloudera Issue: CDH-49165

Routable IP address required by ResourceManager

ResourceManager requires routable `host:port` addresses for `yarn.resourcemanager.scheduler.address`, and does not support using the wildcard `0.0.0.0` address.

Bug: None

Cloudera Bug: CDH-6808

Workaround: Set the address, in the form `host:port`, either in the client-side configuration, or on the command line when you submit the job.

Amazon S3 copy may time out

The Amazon S3 filesystem does not support renaming files, and performs a copy operation instead. If the file to be moved is very large, the operation can time out because S3 does not report progress to the TaskTracker during the operation.

Bug: [MAPREDUCE-972](#)

Cloudera Bug: CDH-17955

Workaround: Use `-Dmapred.task.timeout=15000000` to increase the MR task timeout.

Out-of-memory errors may occur with Oracle JDK 1.8

The total JVM memory footprint for JDK8 can be larger than that of JDK7 in some cases. This may result in out-of-memory errors.

Bug: None

Workaround: Increase max default heap size (`-Xmx`). In the case of MapReduce, for example, increase **Reduce Task Maximum Heap Size** in Cloudera Manager (`mapred.reduce.child.java.opts`, or `mapreduce.reduce.java.opts` for YARN) to avoid out-of-memory errors during the shuffle phase.

MapReduce JAR file renamed (CDH 5.4.0)

As of CDH 5.4.0, `hadoop-test.jar` has been renamed to `hadoop-test-mr1.jar`. This JAR file contains the `mrbench`, `TestDFSIO`, and `nnbench` tests.

Bug: None

Cloudera Bug: CDH-26521

Workaround: None.

Jobs in pool with DRF policy will not run if root pool is FAIR

If a child pool using DRF policy has a parent pool using Fairshare policy, jobs submitted to the child pool do not run.

Affected Versions: All CDH 5 versions.

Bug: [YARN-4212](#)

Cloudera Bug: CDH-31358

Workaround: Change parent pool to use DRF.

Jobs with encrypted spills do not recover if the AM goes down

The fix to [CVE-2015-1776](#) leads to not having enough information to recover a job should the Application Master fail. Releases with this security fix cannot tolerate Application Master failures.

Affected Versions: All CDH 5 versions.

Bug: [MAPREDUCE-6638](#)

Cloudera Bug: CDH-37412

Workaround: None. Fix to come in a later release.

Large TeraValidate data sets can fail with MapReduce

In a cluster using MapReduce, TeraValidate fails when run over large TeraGen/TeraSort data sets (1TB and larger) with an `IndexOutOfBoundsException`. Smaller data sets do not show this issue.

Affected Versions: CDH 5.3.7 and lower

Fixed in Versions: CDH 5.3.8 and higher

Bug: [MAPREDUCE-6481](#)

Cloudera Bug: CDH-31871

Workaround:None.

MapReduce job failure and rolling upgrade (CDH 5.6.0)

MapReduce jobs might fail during a rolling upgrade to or from CDH 5.6.0. Cloudera recommends that you avoid doing rolling upgrades to CDH 5.6.0.

Bug: None

Cloudera Bug: CDH-38587

Workaround: Restart failed jobs.

Unsupported Features

The following features are not currently supported:

- **FileSystemRMStateStore:** Cloudera recommends you use `ZKRMStateStore` (ZooKeeper-based implementation) to store the ResourceManager's internal state for recovery on restart or failover. Cloudera does not support the use of `FileSystemRMStateStore` in production.
- **ApplicationTimelineServer (also known as Application History Server):** Cloudera does not support ApplicationTimelineServer v1. ApplicationTimelineServer v2 is under development and Cloudera does not currently support it.
- **Scheduler Reservations:** Scheduler reservations are currently at an experimental stage, and Cloudera does not support their use in production.
- **Scheduler node-labels:** Node-labels are currently experimental with CapacityScheduler. Cloudera does not support their use in production.
- **CapacityScheduler.** This is deprecated and will be removed from CDH in a future version.

MapReduce1

Oozie workflows not recovered after JobTracker failover on a secure cluster

Delegation tokens created by clients (via `JobClient#getDelegationToken()`) do not persist when the JobTracker [fails over](#). This limitation means that Oozie workflows will not be recovered successfully in the event of a failover on a secure cluster.

Bug: None

Cloudera Bug: CDH-8913

Workaround: Re-submit the workflow.

Hadoop Pipes should not be used in secure clusters

Hadoop Pipes should not be used in secure clusters. A shared password used by the framework for parent-child communications in the clear. A malicious user could intercept that password and potentially use it to access private data in a running application.

Bug: None

No JobTracker becomes active if both JobTrackers are migrated to other hosts

If JobTrackers in an High Availability configuration are shut down, migrated to new hosts, then restarted, no JobTracker becomes active. The logs show a `Mismatched address` exception.

Bug: None

Cloudera Bug: CDH-11801

Workaround: After shutting down the JobTrackers on the original hosts, and before starting them on the new hosts, delete the ZooKeeper state using the following command:

```
$ zkCli.sh rmr /hadoop-ha/<logical name>
```

Hadoop Pipes may not be usable in an MRv1 Hadoop installation done through tarballs

Under MRv1, MapReduce's C++ interface, Hadoop Pipes, may not be usable with a Hadoop installation done through tarballs unless you build the C++ code on the operating system you are using.

Bug: None

Cloudera Bug: CDH-7304

Workaround: Build the C++ code on the operating system you are using. The C++ code is present under `src/c++` in the tarball.

Apache HBase Known Issues

Except for version-specific sections, these issues affect each release of CDH.

Corruption of HBase data stored with MOB feature on upgrade from CDH 5 and HDP 2

CDH 5 and HDP 2 supported the HBase MOB feature prior to that feature making it into an upstream release. When the feature was finalized in the upstream community a different serialization format was used for MOB related configuration settings. CDH 6, HDP 3, and some versions of CDP used this newer serialization format.

When a CDH5 or HDP2 cluster is upgraded to CDH 6, HDP 3 or CDP, this difference in serialization formats causes HBase to incorrectly interpret MOB related configurations. A more severe incorrect interpretation results in the HBase cluster treating the MOB feature as disabled. Clients attempting to read values that were stored in the MOB system will get back HBase internals from the MOB system rather than the underlying data. Additionally, snapshot data which is migrated from CDH 5 and HDP 2 clusters into CDH 6, HDP 3, or CDP clusters may also cause MOB data to be unavailable in the same manner.

If clients perform read-modify-write workloads without recognizing that the value is not what they would expect then the cell value will not be fixed even if the MOB feature is turned back on. For example, using the cell level grant command in the hbase-shell will read cells, update the cell level ACLs on each, and then write the updated cell back to HBase; doing this grant while HBase believes the MOB features is off for a column family with MOB data corrupts any granted cells.

Symptoms:

- MOB configs in table descriptions and HBase web UIs show binary strings instead of sensible values.
- Reading MOB cells results in unexpected values that look like a int followed by a MOB filename.
- If the `IS_MOB` configuration is fixed but `MOB_THRESHOLD` is not, then regions with such configuration settings will fail to open with a `NumberFormatException`.

Products affected:

- CDH
- HDP

Releases affected:

- CDH 5.x
- HDP 2.x

Users affected: Those who use the MOB feature in the affected releases and are planning an upgrade to CDH 6, HDP 3, or CDP Private Cloud.

Severity (Low/Medium/High): High

Impact: Potential data corruption risk for upgrading customers.

Immediate action required:

- **As part of upgrading a cluster that makes use of the MOB feature you must:**
 1. Prior to shutting down the cluster for upgrade, disable any tables which use the MOB feature (look for table descriptions that have column families with `IS_MOB => true`.)

2. Follow normal upgrade steps.
3. Describe any tables which use the MOB feature to confirm that the configuration settings are not correctly displayed.
4. Alter any tables and set the needed MOB configuration settings anew.
5. Describe the tables to confirm the settings show up correctly.
6. Enable the tables.

- **Snapshots**

Snapshots created on MOB-enabled tables prior to upgrade or exported from a CDH 5 and HDP 2 will also result in incorrect `MOB_THRESHOLD` value in the upgraded system.

1. Export a snapshot to a remote filesystem location.
2. Contact Cloudera Support to obtain a tool and instructions to correct the snapshot, *prior* to importing the snapshot into the CDH 6, HDP 3 or CDP HBase filesystem root directory.
3. Use the `clone_snapshot` or the `restore_snapshot` command in HBase on a snapshot originally created in CDH 5 and HDP 2.

- **Repair**

If you have updated a cluster with the MOB feature and data was corrupted prior to correcting the MOB configuration, you can perform a repair using the HFiles still stored in the MOB area of HDFS:

1. Stop all client access.
2. Disable table.
3. Move the table's `mobdir` to `tmp` directory somewhere else in hdfs:

```
hdfs dfs -mv /hbase/mobdir/data/default/<tablename>/<mob region>/
/hbase/.tmp/reload-<table>/
```

4. Enable table.
5. Bulkload `tmp mobdir` from #3:

```
hbase org.apache.hadoop.hbase.tool.LoadIncrementalHFiles /hbase/.tmp/reload/<mob
region> <table_name>
```

6. Verify data.

- **Things to note:**

- The bulkload should be quick as it is moving the file from the `tmp` location into the table data location.
- This process only repair if cell updates did not change timestamp, if it did, you need to delete the cell prior to bulkload.
- As mentioned, bulk loading moves files, it is suggested that you make additional copies rather than just a move to `tmpdir`. This requires more space on HDFS during the process.
- Restored data will not be `mobdir` until compaction occurs. You can run normal compaction or wait until it happens on its own.
- ACLs updates made between upgrade and running the alter for `IS_MOB/MOB_THRESHOLD` will be lost.

Knowledge article: For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-465: Corruption of HBase data stored with MOB feature on upgrade from CDH 5 and HDP 2](#)

HBase suffers data loss during system recovery when a custom WAL directory is configured

[HBASE-20723](#) covers a critical data loss bug. It is present when an HBase deployment is configured to use a non-default location for storing its write-ahead-log. If `hbase.wal.dir` is set to a different location than `hbase.rootdir`, then the recovery process will mistakenly believe there are no edits to replay in the event of process failure of a region server.

Products affected: HBase

Releases affected:

- CDH 5.11.x-5.14.x
- CDH 5.15.0, 5.15.1
- CDH 6.0.0

User affected: Anyone setting the configuration value `hbase.wal.dir` to a setting other than the default. For Cloudera Manager users, this would require setting a safety valve for the `hbase-site.xml` file.

User with non-default setting can determine if they are affected by looking for INFO log messages that indicates edits have been skipped. The following is an example of such message:

```
2018-06-12 22:08:40,455 INFO [RS_LOG_REPLAY_OPS-wn2-duohba:16020-0-Writer-1]
wal.WALSplitter: This region's directory doesn't exist:
hdfs://mycluster/walontest/data/default/tb1/b7fd7db5694eb71190955292b3ff7648. It is very
likely that it was already split so it's safe to discard those edits.
```

Note that the above message is normally harmless, but in this specific edge case the recovery code is looking at the incorrect location to determine region status.

Severity (Low/Medium/High): High

Impact: Data loss is unrecoverable once write-ahead-logs have been deleted as part of routine system processes. There is one exception, if the cluster uses data center replication to ship edits to another cluster and that cluster had not experienced similar data loss.

Immediate action required: Upgrade to a CDH version with the fix.

Addressed in release/refresh/patch: CDH 5.15.2 and higher, CDH 5.16.1 and higher; CDH 6.0.1 and higher

Knowledge article: For the latest update on this issue see the corresponding Knowledge article - [TSB 2019-320: HBase suffers data loss during system recovery when a custom WAL directory is configured](#)

Potential privilege escalation for user of HBase “Thrift 1” API Server over HTTP CVE-2018-8025

CVE-2018-8025 describes an issue in Apache HBase that affects the optional “Thrift 1” API server when running over HTTP. There is a race-condition that could lead to authenticated sessions being incorrectly applied to users, e.g. one authenticated user would be considered a different user or an unauthenticated user would be treated as an authenticated user.

Products affected: HBase Thrift Server

Releases affected:

- CDH 5.4.x - 5.12.x
- CDH 5.13.0, 5.13.1, 5.13.2, 5.13.3
- CDH 5.14.0, 5.14.2, 5.14.3
- CDH 5.15.0

Fixed versions:

- CDH 5.14.4
- CDH 5.15.1

Users affected: Users with the HBase Thrift 1 service role installed and configured to work in “thrift over HTTP” mode. For example, those using Hue with HBase impersonation enabled.

Severity: High

Potential privilege escalation.

CVE: CVE-2018-8025

Immediate action required: Upgrade to a CDH version with the fix, or, disable the HBase Thrift-over-HTTP service. Disabling the HBase Thrift-over-HTTP service will render Hue impersonation inoperable and all HBase access via Hue will be performed using the “hue” user instead of the authenticated user.

Knowledge article: For the latest update on this issue see the corresponding Knowledge article - TSB: 2018-315: Potential privilege escalation for user of HBase “Thrift 1” API Server over HTTP

Known Issues In CDH 5.12.0

IOException from Timeouts

CDH 5.12.0 includes the fix [HBASE-16604](#), where the internal scanner that retries in case of IOException from timeouts could potentially miss data. Java clients were properly updated to account for the new behavior, but thrift clients will now see exceptions where the previous missing data would be.

Workaround: Create a new scanner and retry the operation when encountering this issue.

Known Issues In CDH 5.7.0

Unsupported Features of Apache HBase 1.2

- Although Apache HBase 1.2 allows replication of `hbase:meta`, this feature is not supported by Cloudera and should not be used on CDH clusters until further notice.
- The [FIFO compaction policy](#) has not been thoroughly tested and is not supported in CDH 5.7.0.
- Although Apache HBase 1.2 adds a new permissive mode to allow mixed secure and insecure clients, this feature is not supported by Cloudera and should not be used on CDH clusters until further notice.

The `ReplicationCleaner` process can abort if its connection to ZooKeeper is inconsistent.

Bug: [HBASE-15234](#)

If the connection with ZooKeeper is inconsistent, the `ReplicationCleaner` may abort, and the following event will be logged by the HMaster:

```
WARN org.apache.hadoop.hbase.replication.master.ReplicationLogCleaner: Aborting
ReplicationLogCleaner
because Failed to get list of replicators
```

Unprocessed WALs will accumulate.

Fixed in Versions: CDH 5.3.10 , 5.4.11, 5.5.4, 5.6.1, 5.7.1, 5.8.0, 5.9.0 and above.

Workaround: Restart the HMaster occasionally. The `ReplicationCleaner` will restart if necessary and process the unprocessed WALs.

`IntegrationTestReplication` fails if replication does not finish before the verify phase begins.

Bug: None.

During `IntegrationTestReplication`, if the `verify` phase starts before the `replication` phase finishes, the test will fail because the target cluster does not contain all of the data. If the HBase services in the target cluster does not have enough memory, long garbage-collection pauses might occur

Workaround: Use the `-t` flag to set the timeout value before starting verification.

Cloudera has tested the performance impact of using HDFS encryption with HBase. The overall overhead of HDFS encryption on HBase performance is in the range of 3 to 4% for both read and update workloads. Scan performance has not been thoroughly tested.

Known Issues in CDH 5.6.0

The `ReplicationCleaner` process can abort if its connection to ZooKeeper is inconsistent.

Bug: [HBASE-15234](#)

If the connection with ZooKeeper is inconsistent, the `ReplicationCleaner` may abort, and the following event will be logged by the HMaster:

```
WARN org.apache.hadoop.hbase.replication.master.ReplicationLogCleaner: Aborting
ReplicationLogCleaner
because Failed to get list of replicators
```

Unprocessed WALs will accumulate.

Fixed in Versions: CDH 5.3.10, 5.4.11, 5.5.4, 5.6.1, 5.7.1, 5.8.0, 5.9.0 and above.

Workaround: Restart the HMaster occasionally. The `ReplicationCleaner` will restart if necessary and process the unprocessed WALs.

`ExportSnapshot` or `DistCp` operations may fail on the Amazon `s3a://` protocol.

Bug: None.

`ExportSnapshot` or `DistCP` operations may fail on on AWS when using certain JDK 8 versions, due to an incompatibility between the AWS Java SDK 1.9.x and the `joda-time` date-parsing module.

Workaround: Use `joda-time` 2.8.1 or higher, which is included in AWS Java SDK 1.10.1 or higher.

Reverse scans do not work when Bloom blocks or leaf-level inode blocks are present.

Bug: [HBASE-14283](#)

Because the `seekBefore()` method calculates the size of the previous data block by assuming that data blocks are contiguous, and HFile v2 and higher store Bloom blocks and leaf-level inode blocks with the data, reverse scans do not work when Bloom blocks or leaf-level inode blocks are present when HFile v2 or higher is used.

Fixed in Versions: CDH 5.3.9, 5.4.9, 5.5.2, 5.6.0, 5.7.0, 5.8.0, 5.9.0 and above.

Workaround: None.

Known Issues In CDH 5.5.1

The `ReplicationCleaner` process can abort if its connection to ZooKeeper is inconsistent.

Bug: [HBASE-15234](#)

If the connection with ZooKeeper is inconsistent, the `ReplicationCleaner` may abort, and the following event will be logged by the HMaster:

```
WARN org.apache.hadoop.hbase.replication.master.ReplicationLogCleaner: Aborting
ReplicationLogCleaner
because Failed to get list of replicators
```

Unprocessed WALs will accumulate.

Fixed in Versions: CDH 5.3.10, 5.4.11, 5.5.4, 5.6.1, 5.7.1, 5.8.0, 5.9.0 and above.

Workaround: Restart the HMaster occasionally. The `ReplicationCleaner` will restart if necessary and process the unprocessed WALs.

Extra steps must be taken when upgrading from CDH 4.x to CDH 5.5.1.

The fix for TSB 2015-98 disables legacy object serialization. This will cause direct upgrades on HBase clusters from CDH 4.x to CDH 5.5.1 to fail if one of the workarounds below is not used.

Bug: [HBASE-14799](#)

Cloudera Bug: CDH-34565

Fixed in Versions: CDH 5.3.9, 5.4.9, 5.5.1, 5.6.0, 5.7.0, 5.8.0, 5.9.0 and above.

Workaround: Use one of the following workarounds to upgrade from CDH 4.x to CDH 5.1:

- Upgrade to a CDH 5 version prior to CDH 5.5.1, and then upgrade from that version to CDH 5.5.1, or
- Set the `hbase.allow.legacy.object.serialization` to `true` in the Advanced Configuration Snippet for `hbase-site.xml` if using Cloudera Manager, or directly in `hbase-site.xml` on an unmanaged cluster. Upgrade your cluster to CDH 5.5.1. Remove the `hbase.allow.legacy.object.serialization` property or set it to `false` after migration is complete.

Known Issues In CDH 5.5.0

An operating-system level tuning issue in RHEL7 causes significant latency regressions

There are two distinct causes for the regressions, depending on the workload:

- For a cached workload, the regression may be up to 11%, as compared to RHEL6. The cause relates to differences in the CPU's C-state (power saving state) behavior. With the same workload, the CPU is around 40% busier in RHEL7, and the CPU spends more time transitioning between C-states in RHEL7. Transitions out of deeper C-states add latency. When CPUs are configured to never enter a C-state lower than 1, RHEL7 is slightly faster than RHEL6 on the cached workload. The root cause is still under investigation and may be hardware-dependent.
- For an IO-bound workload, the regression may be up to 8%, even with common C-state settings. A 6% difference in average disk service time has been observed, which in turn seems to be caused by a 10% higher average read size at the drive on RHEL7. The read sizes issued by HBase are the same in both cases, so the root cause seems to be a change in the EXT4 filesystem or the Linux block IO later. The root cause is still under investigation.

Bug: None

Severity: Medium

Workaround: Avoid using RHEL 7 if you have a latency-critical workload. For a cached workload, consider tuning the C-state (power-saving) behavior of your CPUs.

A RegionServer under extreme duress due to back-to-back garbage collection combined with heavy load on HDFS can lock up while attempting to append to the WAL.

The RegionServer appears operational to ZooKeeper, and continues to host regions, but cannot complete any writes. The most obvious symptom is that all writes to regions on this RegionServer time out, and the RegionServer log shows no progress other than queuing of flushes that never complete. Log messages such as the following may occur:

```
124028 2015-11-14 05:54:48,659 WARN org.apache.hadoop.hbase.util.Sleeper: We slept
42911ms instead of 3000ms,
  this is likely due to a long garbage collecting pause and it's usually bad, see
http://hbase.#
124029 2015-11-14 05:54:48,659 WARN org.apache.hadoop.hbase.util.JvmPauseMonitor: Detected
pause in JVM or
  host machine (eg GC): pause of approximately 41110ms
```

```
1806 2015-11-14 04:58:09,952 INFO org.apache.hadoop.hbase.regionserver.wal.FSHLog:
Slow sync cost: 2734 ms, current pipeline:
[DatanodeInfoWithStorage[10.17.198.17:20002,DS-56e2cf88-f267-43a8-b964-b29858#
1807 2015-11-14 04:58:09,952 INFO org.apache.hadoop.hbase.regionserver.wal.FSHLog:
Slow sync cost: 2963 ms, current pipeline:
[DatanodeInfoWithStorage[10.17.198.17:20002,DS-56e2cf88-f267-43a8-b964-b29858#
```

Bug: [HBASE-14374](#)

Workaround: Restart the RegionServer. To avoid the problem, adjust garbage-collection settings, give the RegionServer more RAM, and reduce the load on HDFS.

Known Issues In CDH 5.4 and Higher

The `ReplicationCleaner` process can abort if its connection to ZooKeeper is inconsistent.

Bug: [HBASE-15234](#)

If the connection with ZooKeeper is inconsistent, the `ReplicationCleaner` may abort, and the following event will be logged by the HMaster:

```
WARN org.apache.hadoop.hbase.replication.master.ReplicationLogCleaner: Aborting
ReplicationLogCleaner
because Failed to get list of replicators
```

Unprocessed WALs will accumulate.

Fixed in Versions: CDH 5.3.10, 5.4.11, 5.5.4, 5.6.1, 5.7.1, 5.8.0, 5.9.0 and above.

Workaround: Restart the HMaster occasionally. The `ReplicationCleaner` will restart if necessary and process the unprocessed WALs.

Increments and `CheckAnd*` operations are much slower in CDH 5.4 and higher (since HBase 1.0.0) than in CDH 5.3 and earlier.

This is due to the unification of `mvcc` and `sequenceid` done in HBASE-8763.

Bug: [HBASE-14460](#)

Workaround: None

Known Issues In CDH 5.3 and Higher

Export to Azure Blob Storage (the `wasb://` or `wasbs://` protocol) is not supported.

Bug: [HADOOP-12717](#)

CDH 5.3 and higher supports Azure Blob Storage for some applications. However, a null pointer exception occurs when you specify a `wasb://` or `wasbs://` location in the `--copy-to` option of the `ExportSnapshot` command or as the output directory (the second positional argument) of the `Export` command.

Workaround: None.

Silent Data Loss in Apache HBase Replication

In deployments with slow or unreliable network links, HBase's cross-datacenter replication code may believe it has reached the end of a write-ahead-log prior to successfully parsing the entire file. The remaining content is ignored and never sent to destination clusters. Eventually, the write-ahead-logs are cleaned up as a normal part of HBase operations, preventing later manual recovery.

The issue is fixed by HBASE-15984, which detects when a parsing error has occurred prior to reaching the end of a write-ahead-log file and subsequently retries.

Products affected: HBase and Search

Releases affected: All CDH 5 releases prior to CDH 5.8.0

Users affected: Clusters where HBase replication is enabled.

Severity (Low/Medium/High): High

Impact: The destination cluster will fail to receive all updates, causing inconsistencies between clusters replicating data.

Immediate action required: Customers relying on HBase data center replication should upgrade to CDH 5.8.0 or higher.

Fixed in Versions: CDH 5.8.0 and higher.

HBase in CDH 5 Dependent on Protobuf 2.5

HBase has a transitive dependency on protobuf 2.5. Applications that process protobuf 3 messages generate an error:

```
NoClassDefFoundError: Could not initialize class org.apache.hadoop.hbase.util.ByteStringer
```

Workaround: In CDH 5.8 and higher, use the Apache Maven Shade Plugin to rename protobuf 3.0 packages in the byte code. The Java code looks the same and uses the original package name. However, the byte code contains a different name, so when the HBase client classes load protobuf 2.5, there are no conflicting classes.

Some HBase Features Not Supported in CDH 5

The following features, introduced upstream in HBase, are not supported in CDH 5:

- Visibility labels
- Transparent server-side encryption
- Stripe compaction
- Distributed log replay

Medium Object Blob (MOB) Data Loss of a Snapshot After MOB Compaction

When taking a snapshot of a Medium Object Blobs (MOBs)-enabled table, a race condition can cause data to be lost after a MOB compaction operation completes.

Normally, a table snapshot flushes each region in the table and builds a snapshot manifest with metadata pointing to all the contents of the snapshot. These snapshot manifests are later used to determine which data files can be safely removed after a compaction process. When taking a snapshot, each region's flush can happen in parallel since each region is independent of all the other regions.

When the MOB feature is enabled, there is a separate special MOB region that stores the MOB items. Prior to applying the referenced patch for HBASE-16841, the MOB region was treated as if it was also independent of the other regions. This was not the case – when regions were flushed in a MOB-enabled table, new MOB data may have been written into the special MOB region. Because new MOB data files may have been written out after the MOB region's contents were added to the manifest, some of the MOB data may not have been recorded in the snapshot manifest.

Snapshots on MOB enabled tables were not susceptible to data loss until after a MOB compaction is run. However, after a MOB compaction was run, the MOB data files that were not captured by the snapshot manifest may no longer have had references and were thus eligible for deletion and susceptible to data loss.

This issue is fixed by HBASE-16841.

Releases affected:

CDH 5.3.0, 5.3.1, 5.3.2, 5.3.3, 5.3.4, 5.3.5, 5.3.7, 5.3.8, 5.3.9, 5.3.10

CDH 5.4.0, 5.4.1, 5.4.2, 5.4.3, 5.4.4, 5.4.5, 5.4.7, 5.4.8, 5.4.9, 5.4.10, 5.4.11

CDH 5.5.0, 5.5.1, 5.5.2, 5.5.4, 5.5.5, 5.5.6

CDH 5.6.0, 5.6.1

CDH 5.7.0, 5.7.1, 5.7.2, 5.7.3, 5.7.4, 5.7.5

CDH 5.8.0, 5.8.1, 5.8.2, 5.8.3

CDH 5.9.0, 5.9.1

CDH 5.10.0

Users affected: Only HBase users who enable MOB and take snapshots. Snapshots on regular tables are not affected.

Severity (Low/Medium/High): High

Impact: MOB data may not be properly restored from a HBase Snapshot.

Immediate action required: Upgrade CDH to version 5.7.6, 5.8.4, 5.9.2, 5.10.1 or 5.11.0 and higher

Addressed in release/refresh/patch: Fixed in CDH 5.7.6, 5.8.4, 5.9.2, 5.10.1 or 5.11.0 and higher

Medium Object Blob (MOB) Data Loss After MOB Compaction

If you enable Medium Object Blobs (MOBs) on a table, data loss can occur after a MOB compaction.

When there are no outstanding scanners for HBase regions, by way of optimization, HBase drops cell sequence IDs during normal region compaction for MOB-enabled tables. If a file with no sequence IDs is compacted with an older file that has overlapping cells, the wrong cells may be returned on subsequent compactions. The result is incorrect MOB file references.

The problem manifests as an inability to Scan or Get values from these overlapping rows, and the following WARN-level messages appear in RegionServer logs:

```
WARN HStore Fail to read the cell, the mob file <file name> doesn't exist
java.io.FileNotFoundException: File does not exist:
```

This issue is fixed by HBASE-13922.

Releases affected:

CDH 5.3.0, 5.3.1, 5.3.2, 5.3.3, 5.3.4, 5.3.5, 5.3.7, 5.3.8, 5.3.9, 5.3.10

CDH 5.4.0, 5.4.1, 5.4.2, 5.4.3, 5.4.4, 5.4.5, 5.4.7, 5.4.8, 5.4.9, 5.4.10, 5.4.11

CDH 5.5.0, 5.5.1, 5.5.2, 5.5.4, 5.5.5

CDH 5.6.0, 5.6.1

Users affected: HBase users who enable MOB

Severity (Low/Medium/High): High

Impact: MOB data cannot be retrieved from HBase tables.

Immediate action required: Upgrade CDH to version 5.5.6 or 5.7.0 and higher

UnknownScannerException [Messages After Upgrade](#)

HBase clients may throw exceptions like the following after an HBase upgrade:

```
org.apache.hadoop.hbase.UnknownScannerException:
org.apache.hadoop.hbase.UnknownScannerException: Name: 10092964, already closed?
```

In this upgrade scenario, these messages are caused by restarting the RegionServer during the upgrade. Restart the HBase client to stop seeing the exceptions. The log message has been improved in CDH 5.8.0 and higher.

[HBase moves to Protoc 2.5.0](#)

This change may cause JAR conflicts with applications that have older versions of `protobuf` in their Java classpath.

Bug: None

Workaround: Update applications to use Protoc 2.5.0.

[HBase may not tolerate HDFS root directory changes](#)

While HBase is running, do not stop the HDFS instance running under it and restart it again with a different root directory for HBase.

Bug: None

Cloudera Bug: CDH-5697

Workaround: None

[AccessController postOperation problems in asynchronous operations](#)

When security and Access Control are enabled, the following problems occur:

- If a `Delete Table` fails for a reason other than missing permissions, the access rights are removed but the table may still exist and may be used again.

- If `hbaseAdmin.modifyTable()` is used to delete column families, the rights are not removed from the Access Control List (ACL) table. The `postOperation` is implemented only for `postDeleteColumn()`.
- If `Create Table` fails, full rights for that table persist for the user who attempted to create it. If another user later succeeds in creating the table, the user who made the failed attempt still has the full rights.

Bug: [HBASE-6992](#)

Cloudera Bug: CDH-8566

Workaround: None

Native library not included in tarballs

The native library that enables RegionServer page pinning on Linux is not included in tarballs. This could impair performance if you install HBase from tarballs.

Bug: None

Cloudera Bug: CDH-7304c

Workaround: Use parcels

Apache Hive Known Issues

This topic also contains:

- [HCatalog Known Issues](#)
- [Hive-on-Spark Known Issues](#)

CREATE TABLE EXTERNAL commands run very slowly on Amazon S3

When you run a `CREATE EXTERNAL TABLE` command on Amazon S3 storage, the command runs extremely slowly and might not complete. This usually occurs for tables that have pre-existing, nested data. There is no known workaround.

Affected Versions: CDH 5.15.0

Cloudera Bug: CDH-68833

Workaround: None

Addressed in release/refresh/patch: CDH 5.15.1 and later

Apache Hive vulnerabilities CVE-2018-1282 and CVE-2018-1284

This security bulletin covers two vulnerabilities discovered in Hive.

CVE-2018-1282: JDBC driver is susceptible to SQL injection attack if the input parameters are not properly cleaned

This vulnerability allows carefully crafted arguments to be used to bypass the argument-escaping and clean-up that the Apache Hive JDBC driver does with `PreparedStatement` objects.

If you use Hive in CDH, you have the option of using the Apache Hive JDBC driver or the Cloudera Hive JDBC driver, which is distributed by Cloudera for use with your JDBC applications. Cloudera strongly recommends that you use the Cloudera Hive JDBC driver and offers only limited support for the Apache Hive JDBC driver. If you use the Cloudera Hive JDBC driver you are *not* affected by this vulnerability.

Mitigation: Upgrade to use the [Cloudera Hive JDBC driver](#), or perform the following actions in your Apache Hive JDBC client code or application when dealing with user provided input in `PreparedStatement` objects:

- Avoid passing user input with the `PreparedStatement.setBinaryStream` method, and
- Sanitize user input for the `PreparedStatement.setString` method by replacing all occurrences of `\'` (backslash, single quotation mark) to `'` (single quotation mark).

Detected by: CVE-2018-1282 was detected by Bear Giles of Snaplogic

CVE-2018-1284: Hive UDF series `UDFXPathXXXX` allows users to pass carefully crafted XML to access arbitrary files

If Hive impersonation is disabled and / or Apache Sentry is used, a malicious user might use any of the Hive `xpath` UDFs to expose the contents of a file on the node that is running HiveServer2 which is owned by the HiveServer2 user (usually `hive`).

Mitigation: Upgrade to a release where this is fixed. If `xpath` functions are not currently used, disable them with Cloudera Manager by setting the `hive.server2.builtin.udf.blacklist` property to `xpath,xpath_short,xpath_int,xpath_long,xpath_float,xpath_double,xpath_number,xpath_string` in the **HiveServer2 Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`**. For more information about setting this property to blacklist Hive UDFs, see [the Cloudera Documentation](#).

Products affected: Hive

Releases affected:

- CDH 5.12 and earlier
- CDH 5.13.0, 5.13.1, 5.13.2, 5.13.3
- CDH 5.14.0, 5.14.1, 5.14.2

Users affected: All

Severity (Low/Medium/High): High

Impact: SQL injection, compromise of the `hive` user account

CVE: CVE-2018-1282, CVE-2018-1284

Immediate action required: Upgrade to a CDH release with the fix or perform the above mitigations.

Addressed in the following releases:

- CDH 5.14.4 and higher
- CDH 5.15.0 and higher

For the latest update on these issues, see the corresponding Knowledge article:

[TSB 2018-299: Hive Vulnerabilities CVE-2018-1282 and CVE-2018-1284](#)

Missing results in Hive, Spark, Pig, Custom MapReduce jobs, and other Java applications when filtering Parquet data written by Impala

Apache Hive and Apache Spark rely on Apache Parquet's `parquet-mr` Java library to perform filtering of Parquet data stored in row groups. Those row groups contain statistics that make the filtering efficient without having to examine every value within the row group.

Recent versions of the `parquet-mr` library contain a bug described in [PARQUET-1217](#). This bug causes filtering to behave incorrectly if only some of the statistics for a row group are written. Starting in CDH 5.13, Apache Impala populates statistics in this way for Parquet files. As a result, Hive and Spark may incorrectly filter Parquet data that is written by Impala.

In CDH 5.13, Impala started writing Parquet's `null_count` metadata field without writing the `min` and `max` fields. This is valid, but it triggers the [PARQUET-1217](#) bug in the predicate push-down code of the Parquet Java library (`parquet-mr`). If the `null_count` field is set to a non-zero value, `parquet-mr` assumes that `min` and `max` are also set and reads them without checking whether they are actually there. If those fields are not set, `parquet-mr` reads their default value instead.

For integer SQL types, the default value is 0, so `parquet-mr` incorrectly assumes that the `min` and `max` values are both 0. This causes the problem when filtering data. Unless the value 0 itself matches the search condition, all row groups are discarded due to the incorrect `min/max` values, which leads to missing results.

Affected Products: The Parquet Java library (`parquet-mr`) and by extension, all Java applications reading Parquet files, including, but not limited to:

- Hive
- Spark

- Pig
- Custom MapReduce jobs

Affected Versions:

- CDH 5.13.0, 5.13.1, 5.13.2, and 5.14.0
- CDS 2.2 Release 2 Powered by Apache Spark and earlier releases on CDH 5.13.0 and later

Who Is Affected: Anyone writing Parquet files with Impala and reading them back with Hive, Spark, or other Java-based components that use the `parquet-mr` libraries for reading Parquet files.

Severity (Low/Medium/High): High

Impact: Parquet files containing null values for integer fields written by Impala produce missing results in Hive, Spark, and other Java applications when filtering by the integer field.

Immediate Action Required:

- **Upgrade**

You should upgrade to one of the fixed maintenance releases mentioned below.

- **Workaround**

This issue can be avoided at the price of performance by disabling predicate push-down optimizations:

- In Hive, use the following `SET` command:

```
SET hive.optimize.ppd = false;
```

- In Spark, disable the following configuration setting:

```
--conf spark.sql.parquet.filterPushdown=false
```

Addressed in the Following Releases:

- CDH 5.13.3 and higher
- CDH 5.14.2 and higher
- CDH 5.15.0 and higher
- CDS 2.3 Release 2 and higher

For the latest update on this issue, see the corresponding Knowledge Base article:

[TSB:2018-300: Missing results in Hive, Spark, Pig, and other Java applications when filtering Parquet data written by Impala](#)

Hive Replications Can Fail Intermittently

Hive replication jobs can fail intermittently during the first step of replication while exporting Hive metadata. When this happens, the following message displays in Cloudera Manager: **The remote command failed with error message: Hive Replication Export Step failed.** The likelihood that this error can occur increases with the length of time it takes to export Hive metadata. If you have a very large Hive deployment (those containing many tables and partitions), and/or environments in which the Hive metastore server is under-resourced, you are more likely to see this issue. However, having many tables and partitions or having an under-resourced Hive metastore server do not cause this issue. The problem is caused by a bug and the only solution is to upgrade to a later release. Cloudera strongly recommends upgrading to anyone running Hive Replication and using Cloudera Manager 5.13.0 or 5.13.1 versions.

Affected Products: Cloudera Manager

Affected Versions: Cloudera Manager 5.13.0, 5.13.1

Who is Affected: Anyone using Hive Replication

Severity (Low/Medium/High): High

Cause and Impact: A file is created during the Hive Export Phase, and is used by later phases, such as the Transfer/Import phases. Because of a bug, sometimes the file is overwritten by another process. Hive Replication can thus fail intermittently because of `FileNotFoundException`s. A retry may or may not resolve the issue. This issue can potentially affect some or all Hive replication schedules.

Workaround: Upgrade to Cloudera Manager 5.13.2, 5.14.1, or higher versions

Addressed in release/refresh/patch: Cloudera Manager 5.13.2, 5.14.1, or higher versions.

For the latest updates on this issue, see the following corresponding Knowledge Base article:

[TSB 2018-276: Hive Replications can fail intermittently in Cloudera Manager 5.13.0, 5.13.1 versions](#)

Slow or stalled queries under highly concurrent write workloads when Sentry + Hive are used, caused by a Hive metastore deadlock

In the affected releases, some workloads can cause a deadlock in the Hive metastore. There are internal automated mechanisms to recover from such deadlocks. However, in highly concurrent write-heavy workloads, recovery from the deadlock is slower than the buildup of pending queries or operations. Therefore, the Hive metastore spends a lot of time in the deadlocked state, leading to slow performance or hangs. This in turn impacts performance of HiveServer2 and query performance.

After upgrading to an affected release, if performance of your workloads has worsened dramatically or workloads have stalled, you might be encountering this issue. If MySQL or MariaDB are configured as the Hive metastore database, the following error message appears in the Hive metastore log file, which confirms this issue as the cause:

```
com.mysql.jdbc.exceptions.jdbc4.MySQLTransactionRollbackException: Lock wait
timeout exceeded; try restarting transaction
```



Note: On Oracle and PostgreSQL databases, this message is not printed to the log file, though the issue is database-agnostic.

Affected Versions: CDH 5.13.0, 5.13.1, 5.13.2, 5.14.0, and 5.14.1

Affected Components: When Hive and Sentry are running together

Who is Affected:

- If you are using an affected version of CDH and using Hive and Sentry, refer to the "Workaround" section below.
- If you are using an affected version, but are **not** using Hive **and** Sentry, no action is required.
- If you are not using an affected version and you are using Hive and Sentry, do not upgrade to an affected version of CDH.

Severity (Low/Medium/High): High

Impact: Slow performance of the Hive metastore or Hive queries and eventually the metastore hangs.

Workaround: To mitigate this issue, apply the following configuration changes:

1. In Cloudera Manager, go to your cluster > **Hive** service > **Configuration**, and search for **Hive Metastore Server Advanced Configuration Snippet (Safety Valve) for hive-site.xml**.
2. Click the plus sign (+) twice, and add the following values:
 - Set `hive.metastore.transactional.event.listeners` to an empty value.
 - Set `hive.metastore.event.listeners` to `org.apache.hive.hcatalog.listener.DbNotificationListener`
3. Restart the Hive metastore instances so the configuration changes can be applied.

The following figure shows the configuration details:

Hive Metastore Server Advanced Configuration Snippet (Safety Valve) for hive-site.xml	Hive Metastore Server Default Group	View as XML
Name	<input type="text" value="hive.metastore.transactional.event.listeners"/>	
Value	<input type="text" value="Value"/>	
Description	<input type="text" value="Description"/>	
	<input type="checkbox"/> Final	
Name	<input type="text" value="hive.metastore.event.listeners"/>	
Value	<input type="text" value="org.apache.hive.hcatalog.listener.DbNotificationListener"/>	
Description	<input type="text" value="Description"/>	
	<input type="checkbox"/> Final	

Applying this workaround may cause some DDL queries, such as queries where a table is dropped and a new table is created with the same name, to fail with unexpected `No valid privileges` errors. Re-run such queries to resolve this error.

If the issue persists after applying the above workaround, check the status of the following Knowledge Base article for the latest updates on this issue:

- [TSB 2018-294: Slow or stalled queries under highly concurrent write workloads when Sentry + Hive are used, caused by a Hive metastore deadlock](#)

Addressed in release/refresh/patch: CDH 5.13.3, CDH 5.14.2

Beeline fails with "Error: Unrecognized column type: UNIONTYPE"

When the following conditions are met, Beeline fails and returns the error `Error: Unrecognized column type: UNIONTYPE`:

- Submit a query with Beeline that returns data in the result set from columns that are set to the `UNIONTYPE` data type.
- Beeline option `--incremental=` is set to `true`. By default, this option is set to `false` in CDH.

Affected Version: CDH 5.14 and higher

Cloudera Bug: CDH-63768

Workaround: Set the Beeline option `--incremental=` to `false` on the command line. See [the Beeline documentation](#) for information about setting this option.

DirectSQL with PostgreSQL

Hive doesn't support Hive direct SQL queries with PostgreSQL database. It only supports this feature with MySQL, MariaDB, and Oracle. With PostgreSQL, direct SQL is disabled as a precaution because there have been issues reported upstream. For example, it has been reported that it is not possible to fall back on DataNucleus in case of some failures, plus some other non-standard behaviors. For more information, see [Hive Configuration Properties](#).

Affected Versions: CDH 5.13.0 and higher

Bug: None

Cloudera Bug: CDH-63059

Increase width of columns used for general configuration in the metastore

[HIVE-12274](#) was only partially backported to CDH. The HMS backend database changes are not included in CDH, as it would be a backward compatibility breaking change. The HMS schema won't be changed automatically during the update. It needs to be changed manually.

Affected Versions: CDH 5.13.0 and higher

Bug: [HIVE-12274](#)

Workaround: Change the HMS schema manually.

Hive SSL vulnerability bug disclosure

If you use Cloudera Hive JDBC drivers to connect your applications with HiveServer2, you are *not affected* if:

- SSL is *not* turned on, or
- SSL is turned on but only *non-self-signed* certificates are used.

If neither of the above statements describe your deployment, please read Security Bulletin [Apache Hive SSL Vulnerability Bug Disclosure](#) for further details.

Leak of threads from `getInputPaths` and `getInputSummary` thread pool can degrade HiveServer2 performance

The optimization made to `Utilities.getInputPaths()` method with [HIVE-15546](#) introduced a thread pool that is not shut down when its threads complete their work. This leads to a thread leak for each query that runs in HiveServer2 and the leaked threads are not removed automatically. When queries span multiple partitions, the number of threads spawned increases, but they are never reduced. After 10,000 threads are reached, HiveServer2 performance slows.

Affected Versions: CDH 5.11.0, 5.11.1, 5.11.2, 5.12.0, 5.12.1

Fixed in Versions: 5.12.2, 5.13.0 and higher

Bug: [HIVE-16949](#)

Cloudera Bug: CDH-57789

Resolution: Upgrade to a fixed version of CDH or use the workaround.

Workaround: Upgrading to a fixed version of CDH is the preferred resolution for this issue, but if that is not immediately possible, you can set the **Hive > Configuration > HiveServer2 > Performance > Input Listing Max Threads** property to 1 in Cloudera Manager for managed clusters. If your cluster is not managed, set `hive.exec.input.listing.max.threads=1` in the `hive-site.xml` file.



Important: Decreasing the value set for **Input Listing Max Threads** in Cloudera Manager or decreasing it for the `hive.exec.input.listing.max.threads` property in the `hive-site.xml` file might degrade performance of Hive on Amazon S3, but not for Hive on HDFS.

Hive metastore canary fails

The Hive metastore canary can fail in environments where the Hive-Sentry metastore plugin is deployed and where there are over one million Hive tables and partitions. In this situation, Sentry times out while waiting to receive a full snapshot from the Hive metastore.

Affected Versions: CDH 5.9.2, CDH 5.10.1, CDH 5.10.2, and CDH 5.11.1., 5.11.2, and CDH 5.12

Cloudera Bug: CDH-55255

Resolution: Use workaround.

Workaround: If your environment contains over one million Hive tables and partitions, before you upgrade to one of the affected versions, increase the Sentry timeout property for the metastore plugin to at least 10 to 20 minutes (600,000 to 1,200,000 ms). In Cloudera Manager, set the **Hive Service Advanced Configuration Snippet (Safety Valve) for sentry-site.xml** as follows:

- **Name:** `sentry.hdfs.service.client.server.rpc-connection-timeout`

- **Value:** 600000

After setting this safety valve, monitor your environment. If the Hive metastore canary continues to fail, increase the value by four-minute increments until Sentry can receive the full snapshot from the metastore. For example, if you set the safety valve to 600,000 ms and the canary fails, increase it to 840,000 ms. If it still fails, increase the value to 1,080,000 ms, and so on.

Hive Local Mode is not supported in production environments

Cloudera does not currently support the use of Hive Local Mode in production environments. This mode is set with the `hive.exec.mode.local.auto` property. Use this mode for testing purposes only.

Affected Versions: Not applicable.

Bug: None.

Workaround: None.

INSERT INTO overwrites EXTERNAL tables on the local filesystem

When EXTERNAL tables are located on the local filesystem (URIs beginning with `file://`), the INSERT INTO statement overwrites the table data. Defining EXTERNAL tables on the local filesystem is not a well-documented practice so its behavior is not well defined and is subject to change.

Affected Versions: CDH 5.10

Fixed in Version: CDH 5.10.1

Cloudera Bug: CDH-47856

Workaround: Change the table location from the local filesystem to an HDFS location.

Built-in version() function is not supported

Cloudera does not currently support the built-in `version()` function.

Affected Versions: Not applicable.

Cloudera Bug: CDH-40979

Workaround: None.

EXPORT and IMPORT commands fail for tables or partitions with data residing on Amazon S3

The EXPORT and IMPORT commands fail when the data resides on the Amazon S3 filesystem because the default Hive configuration restricts which file systems can be used for these statements.

Bug: None.

Resolution: Use workaround.

Workaround: Add S3 to the list of supported filesystems for EXPORT and IMPORT by setting the following property in the HiveServer2 Advanced Configuration Snippet (Safety Valve) for `hive-site.xml` in Cloudera Manager (select **Hive** service > **Configuration** > **HiveServer2**):

```
<property>
<name>hive.exim.uri.scheme.whitelist</name>
<value>hdfs,pfile,s3a</value>
</property>
```

Hive queries on MapReduce 1 cannot use Amazon S3 when the Cloudera Manager External Account feature is used

Hive queries that read or write data to Amazon S3 and use the Cloudera Manager External Account feature for S3 credential management do not work with MapReduce 1 (MRv1) because it is deprecated on CDH.

Cloudera Bug: CDH-45201

Resolution: Use workaround.

Workaround: Migrate your cluster from MRv1 to MRv2. See [Migrating from MapReduce \(MRv1\) to MapReduce \(MRv2\)](#).

[ALTER PARTITION does not work on Amazon S3 or between S3 and HDFS](#)

Cloudera recommends that you do not use `ALTER PARTITION` on S3 or between S3 and HDFS.

Cloudera Bug: CDH-42420

[Hive cannot drop encrypted databases in cascade if trash is enabled](#)

Affected Versions: CDH 5.7.0, 5.7.1, 5.7.2, 5.7.3, 5.7.4, 5.7.5, 5.7.6, 5.8.0, 5.8.1, 5.8.2, 5.8.3, 5.8.4, 5.8.5, 5.9.0, 5.9.1, 5.9.2, 5.10.0, 5.10.1, 5.11.0, 5.11.1

Fixed in Versions: CDH 5.7.7, 5.8.6, 5.9.3, 5.10.2, 5.11.2, 5.12.0

Bug: [HIVE-11418](#)

Cloudera Bug: CDH-29913

Workaround: Remove each table using the `PURGE` keyword (`DROP TABLE table PURGE`). After all tables are removed, remove the empty database (`DROP DATABASE database`).

[Potential failure of "alter table <schema>.<table> rename to <schema>.<new_table_name>"](#)

When Hive renames a managed table, it always creates the new renamed table directory under its database directory in order to preserve the database/table hierarchy. The renamed table directory is created under the default database.

Considering that the encryption of a filesystem is part of the evolution hardening of a system (where the system and the data contained in it already exist) and a database can be already created without the location set (because it is not strictly required) and the default database is outside the same encryption zone (or in a no-encryption zone) the alter table rename operation fails.

Affected Version: CDH 5.5 only

Cloudera Bug: CDH-63062

Resolution: Use workaround.

Workaround: Use the following statements:

```
CREATE DATABASE database_encrypted_zone LOCATION
'/hdfs/encrypted_path/database_encrypted_zone';
USE database_encrypted_zone;
```

```
CREATE TABLE rename_test_table LOCATION
'hdfs/encrypted_path/database_encrypted_zone/rename_test';
ALTER TABLE rename_test_table RENAME TO test_rename_table;
```

The renamed table is created under the default database.

[Hive upgrade from CDH 5.0.5 fails on Debian 7.0 if a Sentry 5.0.x release is installed](#)

Upgrading Hive from CDH 5.0.5 to CDH 5.4, 5.3 or 5.2 fails with the following error if a Sentry version later than 5.0.4 and earlier than 5.1.0 is installed. You will see an error such as the following:

```
: error processing
/var/cache/apt/archives/hive_0.13.1+cdh5.2.0+221-1.cdh5.2.0.p0.32~precise-cdh5.2.0_all.deb
  (--unpack):  trying to overwrite '/usr/lib/hive/lib/commons-lang-2.6.jar', which
is also
  in package sentry 1.2.0+cdh5.0.5
```

This is because of a conflict involving `commons-lang-2.6.jar`.

Cloudera Bug: CDH-25450

Workaround: Upgrade Sentry first and then upgrade Hive. Upgrading Sentry deletes all the JAR files that Sentry has installed under `/usr/lib/hive/lib` and installs them under `/usr/lib/sentry/lib` instead.

Hive creates an invalid table if you specify more than one partition with `alter table`

Hive (in all known versions from 0.7) allows you to configure multiple partitions with a single `alter table` command, but the configuration it creates is invalid for both Hive and Impala.

Bug: None

Resolution: Use workaround.

Workaround:

Correct results can be obtained by configuring each partition with its own `alter table` command in either Hive or Impala. For example, the following:

```
ALTER TABLE page_view ADD PARTITION (dt='2008-08-08', country='us') location
'/path/to/us/part080808' PARTITION
(dt='2008-08-09', country='us') location '/path/to/us/part080809';
```

should be replaced with:

```
ALTER TABLE page_view ADD PARTITION (dt='2008-08-08', country='us') location
'/path/to/us/part080808';
ALTER TABLE page_view ADD PARTITION (dt='2008-08-09', country='us') location
'/path/to/us/part080809';
```

Commands run against an Oracle backed metastore may fail

Commands run against an Oracle-backed Metastore fail with error:

```
javax.jdo.JDODataStoreException Incompatible data type for column TBLS.VIEW_EXPANDED_TEXT
: was CLOB (datastore),
but type expected was LONGVARCHAR (metadata). Please check that the type in the datastore
and the type specified in the MetaData are consistent.
```

This error may occur if the metastore is run on top of an Oracle database with the configuration property `datanucleus.validateColumns` set to `true`.

Bug: None

Workaround: Set `datanucleus.validateColumns=false` in the `hive-site.xml` configuration file.

Hive Web Interface is not supported

Cloudera no longer supports the Hive Web Interface because of inconsistent upstream maintenance of this project.

Bug: [DISTRO-77](#)

Resolution: Use workaround

Workaround: Use Hue and Beeswax instead of the Hive Web Interface.

Hive might need additional configuration to make it work in a federated HDFS cluster

Hive jobs normally move data from a temporary directory to a warehouse directory during execution. Hive uses `/tmp` as its temporary directory by default, and users usually configure `/user/hive/warehouse/` as the warehouse directory. Under Federated HDFS, `/tmp` and `/user` are configured as ViewFS mount tables, and so the Hive job will

actually try to move data between two ViewFS mount tables. Federated HDFS does not support this, and the job will fail with the following error:

```
Failed with exception Renames across Mount points not supported
```

Cloudera Bug: CDH-9430

Resolution: No software fix planned; use the workaround.

Workaround: Modify `/etc/hive/conf/hive-site.xml` to allow the temporary directory and warehouse directory to use the same ViewFS mount table. For example, if the warehouse directory is `/user/hive/warehouse`, add the following property to `/etc/hive/conf/hive-site.xml` so both directories use the ViewFS mount table for `/user`.

```
<property>
  <name>hive.exec.scratchdir</name>
  <value>/user/${user.name}/tmp</value>
</property>
```

Cannot create archive partitions with external HAR (Hadoop Archive) tables

`ALTER TABLE ... ARCHIVE PARTITION` is not supported on external tables.

Cloudera Bug: CDH-9638

Workaround: None

Setting `hive.optimize.skewjoin` to true causes long running queries to fail

Bug: None

Workaround: None

Object types Server and URI are not supported in "`SHOW GRANT ROLE roleName on OBJECT objectName`"

Cloudera Bug: CDH-19430

Workaround: Use `SHOW GRANT ROLE roleName` to list all privileges granted to the role.

Kerberized HS2 with LDAP authentication fails in a multi-domain LDAP case

In CDH 5.7, Hive introduced a feature to support HS2 with Kerberos plus LDAP authentication; but it broke compatibility with multi-domain LDAP cases on CDH 5.7.x and C5.8.x versions.

Affected Versions: CDH 5.7.1, CDH 5.8.0, and CDH 5.8.1

Fixed in Versions: CDH 5.7.2 and higher, CDH 5.8.2 and higher

Bug: [HIVE-13590](#).

Workaround: None.

With encrypted HDFS, 'drop database if exists `<db_name>` cascade' fails

Hive cannot drop encrypted databases in cascade if trash is enabled.

Affected Versions: CDH 5.7.0 - 5.7.6, CDH 5.8.0 - 5.8.5, CDH 5.9.0 - 5.9.2, CDH 5.10.0 - 5.10.1, CDH 5.11.0 - 5.11.1

Fixed in Versions: CDH 5.7.7, 5.8.6, 5.9.3, 5.10.2, 5.11.2, 5.12.0

Bug: [HIVE-11418](#)

Cloudera Bug: CDH-29913

Workaround: Remove each table, using the PURGE keyword (`DROP TABLE table PURGE`). After all tables are removed, remove the empty database (`DROP DATABASE database`).

HCatalog Known Issues



Note: As of CDH 5, HCatalog is part of Apache Hive.

WebHCatalog does not work in a Kerberos-secured federated cluster

Cloudera Bug: CDH-12416

Resolution: None planned.

Workaround: None

Hive-on-Spark (HoS) Known Issues

Hive on Spark queries are failing with "Timed out waiting for client to connect" for an unknown reason

If this exception is preceded by logs of the form "client.RpcRetryingCaller: Call exception...". Then this failure is due to an unavailable HBase service. On a secure cluster, spark-submit will try to obtain delegation tokens from HBase, even though HoS doesn't necessarily need them. So if HBase is unavailable, then spark-submit throws an exception.

Cloudera Bug: CDH-59591

Affected Versions: CDH 5.7.0 and higher

Workaround: Fix the HBase service, or set `spark.yarn.security.tokens.hbase.enabled` to `false`.

Intermittent inconsistent results for Hive-on-Spark with DECIMAL-type columns



Important: If you use Hive with MapReduce2 (YARN) as the execution engine you are unaffected.

SQL queries executed using Hive-on-Spark (HoS) can produce inconsistent results intermittently if they use a `GROUP BY` clause (user-defined aggregate functions, or UDAFs) on `DECIMAL` data type columns.

Hive queries are affected if all of the following conditions are met:

- Spark is used as the execution engine (`hive.execution.engine=spark`) in the Affected Versions listed below.
- The property `spark.executor.cores` is set to a value greater than 1.
- The query uses a UDAF that requires a `GROUP BY` clause on a `DECIMAL` data type column.

For example:

```
COUNT(DISTINCT<decimal_column_name>), SELECT SUM(<column>) GROUP BY<decimal_column>
```

You can use the `EXPLAIN` command to view the query execution plan to check for the presence of the `GROUP BY` clause on a `DECIMAL` data type column.

Bug: [HIVE-16257](#), [HIVE-12768](#)

Severity: High

Affected Versions:

- CDH 5.4.0, 5.4.1, 5.4.2, 5.4.3, 5.4.4, 5.4.5, 5.4.7, 5.4.8, 5.4.9, 5.4.10, 5.4.11
- CDH 5.5.0, 5.5.1, 5.5.2, 5.5.3, 5.5.4, 5.5.5, 5.5.6
- CDH 5.6.0, 5.6.1
- CDH 5.7.0, 5.7.1, 5.7.3, 5.7.4, 5.7.5, 5.7.6
- CDH 5.8.0, 5.8.1, 5.8.2, 5.8.3, 5.8.4
- CDH 5.9.0, 5.9.1

- CDH 5.10.0, 5.10.1

Resolution: Use the workaround or upgrade to a release that addresses this issue:

- CDH 5.8.5 and higher
- CDH 5.9.2 and higher
- CDH 5.11.0 and higher

Workaround: To fix this issue, set `spark.executor.cores=1` for sessions that run queries that use the `GROUP BY` clause on `DECIMAL` data type columns. This might cause performance degradation for some queries. Cloudera recommends that you run performance validations before setting this property.

You can set this with Cloudera Manager in HiveServer2 at the cluster level:



Warning: Setting `spark.executor.cores` to 1 at the cluster-level might impact unaffected queries in HoS (those that do not use UDAFs using the `GROUP BY` clause). Cloudera recommends that you run performance validations before setting this property.

1. In the Cloudera Manager Admin console, go to the Hive service, and select **Configuration>HiveServer2>Performance**.
2. In the **Search** text box, type `spark.executor.cores` and press Enter.
3. Set **Spark Executor Cores** to 1.
4. Click **Save Changes**.

Hive-on-Spark throws exception for a multi-insert with a join query

A multi-insert combined with a join query with Hive on Spark (Hos) sometimes throws an exception. It occurs only when multiple parts of the resultant operator tree are executed on the same executor by Spark.

Affected Versions: CDH 5.7.0

Fixed in Versions: CDH 5.7.1

Bug: [HIVE-13300](#)

Cloudera Bug: CDH-38458

Workaround: Run inserts one at a time.

NullPointerException thrown when a Spark session is reused to run a mapjoin

Some Hive-on-Spark (HoS) queries might fail with a `NullPointerException` if a Spark dependency is not set.

Affected Version: CDH 5.7.0

Bug: [HIVE-12616](#)

Cloudera Bug: CDH-38514

Workaround: Configure Hive to depend on the Spark (on YARN) service in Cloudera Manager.

Large Hive-on-Spark queries might fail in Spark tasks with `ExecutorLostFailure`

The root cause is `java.lang.OutOfMemoryError: Unable to acquire XX bytes of memory, got 0`. Spark executors can OOM because of a failure to correctly spill shuffle data from memory to disk.

Cloudera Bug: CDH-38642

Workaround: Run this query using MapReduce.

Hive-on-Spark2 is not Supported

Hive-on-Spark is a CDH component that has a dependency on Spark 1.6. Because CDH components do not have any dependencies on Spark 2, Hive-on-Spark does not work with the Cloudera Distribution of Apache Spark 2.

Hue Known Issues

High DDL usage in Hue Impala Editor may issue flood of INVALIDATE Calls

Issuing DDL statements using Hue's Impala editor or invoking Hue's "Refresh Cache" function in the left-side metadata browser results in Hue issuing INVALIDATE METADATA calls to the Impala service. This call is expensive and can result in a significant system impact, up to and including full system outage, when repeated in sufficient volume. This has been corrected in HUE-8882.

Components affected:

- Hue
- Impala

Products affected:

- Cloudera Enterprise 5
- Cloudera Enterprise 6

Releases affected:

- CDH 5.15.1, 5.15.2
- CDH 5.16.x
- CDH 6.1.1
- CDH 6.2.x
- CDH 6.3.0, 6.3.1, 6.3.2, 6.3.3

Users affected: End-users using Impala editor in Hue.

Severity: High

Impact: Users running DDL statements using the Hue Impala editor or invoking Hue's Refresh Cache function causes INVALIDATE METADATA commands to be sent to Impala. Impala's metadata invalidation is an expensive operation and could cause impact on the performance of subsequent queries, hence leading to the potential for significant impact on the entire cluster, including the potential for whole-system outage.

Action required:

- CDH 6.x customers: Upgrade to CDH 6.3.4 that contains the fix.
- CDH 5.x customers: Contact Cloudera Support for further assistance.

Apache Issue: [HUE-8882](#)

Knowledge Article: For the latest update on this issue see the corresponding Knowledge article: [Cloudera Customer Advisory: High DDL usage in Hue Impala Editor may issue flood of INVALIDATE Calls](#)

Hue Silently Disables StartTLS in LDAP Connections

There are two mechanisms to secure communication to an LDAP server. One is to use an 'ldaps' connection, where all traffic is encrypted inside a TLS tunnel - much like 'https'. The other is to use 'StartTLS', where traffic begins unencrypted in the "ldap" protocol and then upgrades itself to a TLS connection.

If StartTLS is enabled in the Hue configuration but the 'ldap_cert' parameter is not configured, then Hue silently disables StartTLS.

StartTLS will not be used for synchronization or import, even if StartTLS is enabled and the 'ldap_cert' parameter is set.

The result is that connections that the administrator assumes to be secured, using StartTLS, are not actually secure.

CVE: CVE-2019-19146

Date/time of detection: 22nd March, 2019

Detected by: Ben Gooley, Cloudera

Severity (Low/Medium/High): 8.8 High [CVSS AV:N/AC:L/PR:N/UI:R/S:U/C:H/I:H/A:H](#)

Products affected: CDH

Releases affected:

- CDH 5.x
- CDH 6.1.0
- CDH 6.1.1
- CDH 6.2.0
- CDH 6.2.1
- CDH 6.3.0

Users affected: All users who are using StartTLS enabled in the Hue configuration when using LDAP as Authentication Backend to login in Hue.

Impact: Sensitive data exposure.

Immediate action required:

- **Upgrade (recommended):** Update to a version of CDH containing the fix.
- **Workaround:** Use “ldaps” instead of “ldap” and StartTLS.

Addressed in release/refresh/patch: CDH 6.3.1 and above

Knowledge Article: For the latest update on this issue see the corresponding Knowledge article: [TSB 2020-371: Hue Silently Disables StartTLS in LDAP Connections](#)

Hue allows unsigned SAML assertions

If Hue receives an unsigned assertion, it continues to process it as valid. This means it is possible for an end-user to forge or remove the signature and manipulate a SAML assertion to gain access without a successful authentication.

Products affected: Hue, CDH

Releases affected:

- CDH 5.15.x and earlier
- CDH 5.16.0, 5.16.1
- CDH 6.0.x
- CDH 6.1.x

User affected: All users who are using SAML with Hue.

CVE: CVE-2019-14775

Date/time of detection: January 2019

Detected by: Joel Snape

Severity (Low/Medium/High): High

Impact:

This is a significant security risk as it allows anyone to fake their access validity and therefore access Hue, even if they should not have access. In more detail: if Hue receives an unsigned assertion, it continues to process it as valid. This means it is possible for an end-user to forge or remove the signature and manipulate a SAML assertion to gain access without a successful authentication.

CVE: CVE-2019-14775

Immediate action required:

- **Upgrade (recommended):** Upgrade to a version of CDH containing the fix.
- **Workaround:** None

Addressed in release/refresh/patch:

- CDH 5.16.2
- CDH 6.2.0

Oozie workflows cannot be clicked when opened from **Workflows > Editor > Workflows** in Hue 3

When you are using Hue 3, you cannot edit an Oozie workflow by clicking it when you navigate to **Workflows > Editor > Workflows**.

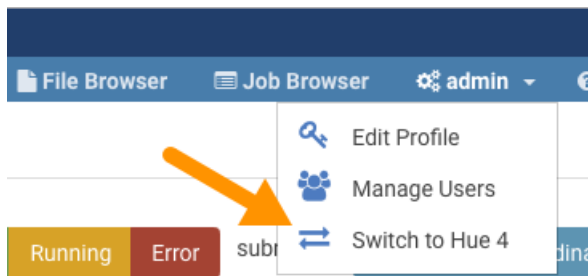
Affected Versions: CDH 5.15.x and later CDH 5.x versions

Fixed in Versions: No fix is planned for CDH 6.x as only Hue 4 is available

Cloudera Bug: CDH-70450

Workaround: Switch to Hue 4:

In the Hue user interface, click the **Administration** menu and select **Switch to Hue 4**:



Long delay in Impala metadata operations cause a 502 proxy error

A 502 error message appears when an Impala metadata operation, such as an Impala query compilation, takes several minutes. This causes Hue to timeout and throw a 502 error.

Affected Versions: CDH 5.12.x, 5.13.0, 5.13.1

Fixed in Versions: 5.14.0

Cloudera Bug: CDH-57588

Workaround: Configure the `ProxyTimeout` value parameter using the Advanced Configuration Snippet:

1. In the Cloudera Manager user interface, locate **Hue > Configuration**.
2. Using the search box, search for 'Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf'.
3. In the resulting text box, insert `ProxyTimeout 600` — this allows the Hue Load Balancer 600 seconds until timeout instead of the default 60.
4. Save and restart the Hue Load Balancer.
5. Export the results to excel again or rerun the Hive query from Hue UI.

Hue operations randomly fail and display a 502 error

Hue operations randomly fail.

Affected Versions: CDH 5.12.x, 5.13.0

Fixed in Versions: 5.14.0

Workaround: Make configuration changes to the Hue Load Balancer:

1. Find the Hue Load Balancer host by selecting **Host > All Host** and search for "Hue Load Balancer"
2. Navigate to the following directory on the Hue Load Balancer host: `/usr/share/cmf/hue/httpd/`.
3. Edit the `httpd.conf` file and add the following property after the `ProxyPass` property: `SetEnv proxy-initial-not-pooled 1`.

4. Inside the `<IfDefine Cloudera_HTTPD_USE_SSL>` block under the `SSLProtocol` property line, add the following properties:

```
SSLProxyProtocol all -SSLv2 -SSLv3
SSLCipherSuite
```

Cloudera Bug: CDH-57588

Hue does not start with MySQL and Ubuntu 12.04

In Ubuntu 12.04, the pip script file, `/usr/lib/python2.7/dist-packages/pkg_resources.py`, is unable to resolve required packages.

Affected Versions: All CDH 5 versions.

Cloudera Bug: CDH-54615

Fixed in Versions: No fix is planned. Recommend upgrading from Ubuntu 12.04

Workaround:

1. Log on to Cloudera Manager and go to **Hue > Configuration**.
2. Search for and set **Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini**:
`PYTHONPATH=/opt/cloudera/parcels/CDH/lib/hue/build/env/lib/python2.7/site-packages`
3. Click **Save Changes** and restart Hue.

Hue Load Balancer with TLS fails on peer Cloudera Navigator validation

The Hue load balancer may fail with TLS enabled and Apache HTTPD 2.4 or higher.

Affected Versions: CDH 5.11.0 and lower

Cloudera Bug: CDH-51352

Workaround:

1. Log on to Cloudera Manager and go to **Hue > Configuration**.
2. Search for and set **Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf**:

```
ProxyPreserveHost Off
```

3. Search for and set **Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini**:

```
[desktop]
use_x_forwarded_host=true
```

4. Click **Save Changes** and restart Hue.

HTTPD 2.4 breaks Hue high availability with SSLProxyPeerCN on

In HTTPD 2.2 and earlier versions, `SSLProxyPeerCN` is off. In HTTPD 2.4, `SSLProxyPeerCN` is turned on. This breaks Hue high availability and throws a proxy error.

Affected Versions: CDH 5.11.0 only

Workaround: Disable `SSLProxyCheckPeerCN` and `SSLProxyCheckPeerName` in `httpd.conf`:

1. Log on to Cloudera Manager and go to **Hue > Configuration**.
2. Search for **Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf**.
3. Disable proxy peer checks by entering:

```
SSLProxyCheckPeerCN off
SSLProxyCheckPeerName off
```

4. Click **Save changes** and restart Hue.

Icons Stop Rendering in IE11 Browser with HTTPS

When you use the back button in IE11 with https, some Hue graphics stop rendering. Add a load balancer to resolve.

Affected Versions: All CDH 5.x versions.

Cloudera Bug: CDH-52950

Workaround: Add a Hue Load Balancer:

1. Log on to Cloudera Manager and click **Hue**.
2. Select **Actions > Add Role Instances**.
3. Click **Select hosts** in the field under **Load Balancer**.
4. Select a host and click **OK**.
5. Check the Load Balancer box and select **Actions for Selected > Start > Start**.
6. Click **Save Changes** and restart Hue.

Hue does not support the Spark App

Hue does not currently support the Spark application.

Apache Impala Known Issues

The following sections describe known issues and workarounds in Impala, as of the current production release. This page summarizes the most serious or frequently encountered issues in the current release, to help you make planning decisions about installing and upgrading. Any workarounds are listed here. The bug links take you to the Impala issues site, where you can see the diagnosis and whether a fix is in the pipeline.



Note: The online issue tracking system for Impala contains comprehensive information and is updated in real time. To verify whether an issue you are experiencing has already been reported, or which release an issue is fixed in, search on the [Impala JIRA tracker](#).

For issues fixed in various Impala releases, see [Fixed Issues in Apache Impala](#) on page 2384.

Impala/Sentry security roles mismatch after Catalog Server restart

This issue occurs when Impala's Catalog Server is restarted without also restarting all the Impala Daemons.

Impala uses generated numeric identifiers for roles. These identifiers are regenerated during catalogd restarts, and the same role can get a different identifier, possibly used by a different role before restart. An impalad's metadata cache can contain old id + role pairs, and when it is updated with privileges with new role ids from the catalog, the privilege will be added to the wrong role; the one that previously had the same role id.

Products affected: Apache Impala

Releases affected:

- CDH 5.14.4 and all prior releases

Users affected: Impala users with authorization enabled.

Date/time of detection: 5th October, 2018

Severity (Low/Medium/High): 3.8 "Low"; [CVSS:3.0/AV:N/AC:H/PR:H/UI:R/S:U/C:L/I:L/A:L](#)

Impact: Users may get privileges of unrelated users.

CVE: CVE-2019-16381

Immediate action required: Update to a version of CDH containing the fix.

Addressed in release/refresh/patch: CDH 5.15.0

Knowledge article: For the latest update on this issue see the corresponding Knowledge article: [TSB 2020-348: Impala/Sentry security roles mismatch after Catalog Server restart](#)

Timestamp type-casted to varchar in a binary predicate can produce incorrect result

In an Impala query the timestamp can be type-casted to a varchar of smaller length to convert a timestamp value to a date string. However, if such Impala query is used in a binary comparison against a string literal, it can produce incorrect results, because of a bug in the expression rewriting code. The following is an example of this:

```
> select * from (select cast('2018-12-11 09:59:37' as timestamp) as ts) tbl where cast(ts as varchar(10)) = '2018-12-11';
```

The output will have 0 rows.

Products affected: Apache Impala

Releases affected:

- CDH 5.15.0, 5.15.1, 5.15.2, 5.16.0, 5.16.1
- CDH 6.0.0, 6.0.1, 6.1.0, 6.1.1

For the latest update on this issue see the corresponding Knowledge article: [TSB 2019-358: Timestamp type-casted to varchar in a binary predicate can produce incorrect result](#)

XSS Cloudera Manager

Malicious Impala queries can result in Cross Site Scripting (XSS) when viewed in Cloudera Manager.

Products affected: Apache Impala

Releases affected:

- Cloudera Manager 5.13.x, 5.14.x, 5.15.1, 5.15.2, 5.16.1
- Cloudera Manager 6.0.0, 6.0.1, 6.1.0

Users affected: All Cloudera Manager Users

Date/time of detection: November 2018

Severity (Low/Medium/High): High

Impact: When a malicious user generates a piece of JavaScript in the impala-shell and then goes to the Queries tab of the Impala service in Cloudera Manager, that piece of JavaScript code gets evaluated, resulting in an XSS.

CVE: CVE-2019-14449

Immediate action required: There is no workaround, upgrade to the latest available maintenance release.

Addressed in release/refresh/patch:

- Cloudera Manager 5.16.2
- Cloudera Manager 6.0.2, 6.1.1, 6.2.0, 6.3.0

Impala Known Issues: Startup

These issues can prevent one or more Impala-related daemons from starting properly.

Problem retrieving FQDN causes startup problem on kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a kerberized cluster.

This problem might occur immediately after an upgrade of a CDH cluster, due to changes in Cloudera Manager that supplies the `--hostname` flag automatically to the Impala-related daemons. (See the issue “hostname parameter is not passed to Impala catalog role” at [the Cloudera Manager Known Issues page](#).)

Bugs: [IMPALA-4978](#), [IMPALA-5253](#)

Severity: High

Resolution: The issue is expected to occur less frequently on systems with fixes for [IMPALA-4978](#), [IMPALA-5253](#), or both. Even on systems with fixes for both of these issues, the workaround might still be required in some cases.

Workaround: Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname` only returns the short name, pass the command-line flag `--hostname=fully_qualified_domain_name` in the startup options of all Impala-related daemons.

Impala Known Issues: Crashes and Hangs

These issues can cause Impala to quit or become unresponsive.

Unable to view large catalog objects in catalogd Web UI

In `catalogd` Web UI, you can list metadata objects and view their details. These details are accessed via a link and printed to a string formatted using thrift's `DebugProtocol`. Printing large objects (> 1 GB) in Web UI can crash `catalogd`.

Bug: [IMPALA-6841](#)

Crash when querying tables with "\0" as a row delimiter

When querying a textfile-based Impala table that uses `\0` as a new line separator, Impala crashes.

The following sequence causes `impalad` to crash:

```
create table tab_separated(id bigint, s string, n int, t timestamp, b boolean)
  row format delimited
  fields terminated by '\t' escaped by '\\' lines terminated by '\000'
  stored as textfile;
select * from tab_separated; -- Done. 0 results.
insert into tab_separated (id, s) values (100, ''); -- Success.
select * from tab_separated; -- 20 second delay before getting "Cancelled due to unreachable
  impalad(s): xxxx:22000"
```

Bug: [IMPALA-6389](#)

Workaround: Use an alternative delimiter, e.g. `\001`.

Resolution: Fixed in Impala 2.13 and higher.

Altering Kudu table schema outside of Impala may result in crash on read

Creating a table in Impala, changing the column schema outside of Impala, and then reading again in Impala may result in a crash. Neither Impala nor the Kudu client validates the schema immediately before reading, so Impala may attempt to dereference pointers that aren't there. This happens if a string column is dropped and then a new, non-string column is added with the old string column's name.

Bug: [IMPALA-4828](#)

Severity: High

Workaround: Run the statement `REFRESH table_name` after any occasion when the table structure, such as the number, names, and data types of columns, are modified outside of Impala using the Kudu API.

Resolution: Fixed in CDH 5.12 / Impala 2.9 and higher.

Queries that take a long time to plan can cause webserver to block other queries

Trying to get the details of a query through the debug web page while the query is planning will block new queries that had not started when the web page was requested. The web UI becomes unresponsive until the planning phase is finished.

Bug: [IMPALA-1972](#)

Severity: High

Resolution: Fixed in CDH 5.12 / Impala 2.9 and higher.

Linking IR UDF module to main module crashes Impala

A UDF compiled as an LLVM module (.ll) could cause a crash when executed.

Bug: [IMPALA-4595](#)

Severity: High

Resolution: Fixed in CDH 5.10 / Impala 2.8 and higher.

Workaround: Compile the external UDFs to a .so library instead of a .ll IR module.

Secure Impala-Kudu clusters need to be restarted frequently

Due to KUDU-2264, secure, long-running Impala clusters on CDH 5.13.x and later will be unable to interact with Kudu tables after some time because the Kudu client used by Impala will no longer be able to authenticate with Kudu. The problem starts after the client's ticket can no longer be renewed, which is typically after 7 days. Restarting the Impala service will allow Impala to work with Kudu again.

Impala will begin failing to read or write Kudu tables with errors in the Impala daemon logs resembling:

```

W1121 22:47:47.231425 50416 ConnectToCluster.java:302] Error receiving response from
cdh-master-4db933ef.cdh-cluster.internal:7051
Java exception follows:
org.apache.kudu.client.NonRecoverableException: Server requires Kerberos, but this client
is not
authenticated (kinit)

at org.apache.kudu.client.Negotiator.evaluateChallenge(Negotiator.java:705)
at org.apache.kudu.client.Negotiator.sendSaslInitiate(Negotiator.java:581)
at org.apache.kudu.client.Negotiator.startAuthentication(Negotiator.java:545)
at org.apache.kudu.client.Negotiator.handleTlsMessage(Negotiator.java:499)
at org.apache.kudu.client.Negotiator.handleResponse(Negotiator.java:264)
at org.apache.kudu.client.Negotiator.messageReceived(Negotiator.java:231)
at ...
Caused by: javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException:
No valid credentials provided (Mechanism level: Failed to find any Kerberos tgt)]

at com.sun.security.sasl.gsskerb.GssKrb5Client.evaluateChallenge(GssKrb5Client.java:211)
at org.apache.kudu.client.Negotiator$1.run(Negotiator.java:691)
at org.apache.kudu.client.Negotiator$1.run(Negotiator.java:688)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:422)
at org.apache.kudu.client.Negotiator.evaluateChallenge(Negotiator.java:687)
... 35 more

Caused by: GSSException: No valid credentials provided (Mechanism level: Failed to find
any
Kerberos tgt)

at sun.security.jgss.krb5.Krb5InitCredential.getInstance(Krb5InitCredential.java:147)
at sun.security.jgss.krb5.Krb5MechFactory.getCredentialElement(Krb5MechFactory.java:122)
at sun.security.jgss.krb5.Krb5MechFactory.getMechanismContext(Krb5MechFactory.java:187)
at sun.security.jgss.GSSManagerImpl.getMechanismContext(GSSManagerImpl.java:224)
at sun.security.jgss.GSSContextImpl.initSecContext(GSSContextImpl.java:212)
at sun.security.jgss.GSSContextImpl.initSecContext(GSSContextImpl.java:179)
at com.sun.security.sasl.gsskerb.GssKrb5Client.evaluateChallenge(GssKrb5Client.java:192)
... 40 more

```

Alternatively, errors may be seen like the following:

```

I0102 11:58:36.688992 3587 jni-util.cc:196]
org.apache.impala.common.ImpalaRuntimeException: Unable to initialize the Kudu scan node

at org.apache.impala.planner.KuduScanNode.init(KuduScanNode.java:133)
at org.apache.impala.planner.SingleNodePlanner.createScanNode(SingleNodePlanner.java:1312)
at
org.apache.impala.planner.SingleNodePlanner.createTableRefNode(SingleNodePlanner.java:1519)
at
org.apache.impala.planner.SingleNodePlanner.createTableRefsPlan(SingleNodePlanner.java:779)
at

```

```

org.apache.impala.planner.SingleNodePlanner.createSelectPlan(SingleNodePlanner.java:617)
at org.apache.impala.planner.SingleNodePlanner.createQueryPlan(SingleNodePlanner.java:260)
at
org.apache.impala.planner.SingleNodePlanner.createSingleNodePlan(SingleNodePlanner.java:150)
at org.apache.impala.planner.Planner.createPlan(Planner.java:98)
at org.apache.impala.service.Frontend.createExecRequest(Frontend.java:1004)
at org.apache.impala.service.Frontend.createExecRequest(Frontend.java:1100)
at org.apache.impala.service.JniFrontend.createExecRequest(JniFrontend.java:156)

Caused by: org.apache.kudu.client.NonRecoverableException: cannot re-acquire
authentication
token after 5 attempts

at
org.apache.kudu.client.AuthnTokenReacquirer$1NewAuthnTokenErrB.failQueuedRpcs(AuthnTokenReacquirer.java:167)
at
org.apache.kudu.client.AuthnTokenReacquirer$1NewAuthnTokenErrB.call(AuthnTokenReacquirer.java:159)
at
org.apache.kudu.client.AuthnTokenReacquirer$1NewAuthnTokenErrB.call(AuthnTokenReacquirer.java:142)
at com.stumbleupon.async.Deferred.doCall(Deferred.java:1262)
at com.stumbleupon.async.Deferred.runCallbacks(Deferred.java:1241)
at com.stumbleupon.async.Deferred.callback(Deferred.java:989)
at
org.apache.kudu.client.ConnectToCluster.incrementCountAndCheckExhausted(ConnectToCluster.java:223)
at org.apache.kudu.client.ConnectToCluster.access$000(ConnectToCluster.java:48)
at
org.apache.kudu.client.ConnectToCluster$ConnectToMasterErrCB.call(ConnectToCluster.java:304)
at
org.apache.kudu.client.ConnectToCluster$ConnectToMasterErrCB.call(ConnectToCluster.java:293)
at com.stumbleupon.async.Deferred.doCall(Deferred.java:1262)
at com.stumbleupon.async.Deferred.runCallbacks(Deferred.java:1241)
at com.stumbleupon.async.Deferred.handleContinuation(Deferred.java:1297)

```

Products Affected: Impala

Releases Affected: CDH 5.13.0, CDH 5.13.1, and CDH 5.14.0

Users Affected: Users running a secure cluster with Impala and Kudu.

Severity: Medium

Impact: Impala will not be able to query or write to Kudu tables until the service is restarted. The Impala service will need to be restarted periodically in order to continue being able to work with Kudu tables.

Immediate action required: To work around the issue, users can restart the Impala service at least once per ticket renewal period. This period is typically 7 days.

Resolution: Addressed in CDH 5.13.2, CDH 5.13.3, and CDH 5.14.2

Setting BATCH_SIZE query option too large can cause a crash

Using a value in the millions for the BATCH_SIZE query option, together with wide rows or large string values in columns, could cause a memory allocation of more than 2 GB resulting in a crash.

Bug: [IMPALA-3069](#)

Severity: High

Resolution: Fixed in CDH 5.9 / Impala 2.7 and higher.

Impala should not crash for invalid avro serialized data

Malformed Avro data, such as out-of-bounds integers or values in the wrong format, could cause a crash when queried.

Bug: [IMPALA-3441](#)

Severity: High

Resolution: Fixed in CDH 5.9.0 / Impala 2.7.0 and CDH 5.8.2 / Impala 2.6.2.

Queries may hang on server-to-server exchange errors

The `DataStreamSender::Channel::CloseInternal()` does not close the channel on an error. This causes the node on the other side of the channel to wait indefinitely, causing a hang.

Bug: [IMPALA-2592](#)

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0.

Impalad is crashing if udf jar is not available in hdfs location for first time

If the JAR file corresponding to a Java UDF is removed from HDFS after the Impala `CREATE FUNCTION` statement is issued, the `impalad` daemon crashes.

Bug: [IMPALA-2365](#)

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0.

Impala Known Issues: Performance

These issues involve the performance of operations such as queries or DDL statements.

Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

Bug: [IMPALA-6671](#)

Profile timers not updated during long-running sort

If you have a query plan with a long-running sort operation, e.g. minutes, the profile timers are not updated to reflect the time spent in the sort until the sort starts returning rows.

Bug: [IMPALA-5200](#)

Workaround: Slow sorts can be identified by looking at "Peak Mem" in the summary or "PeakMemoryUsage" in the profile. If a sort is consuming multiple GB of memory per host, it will likely spend a significant amount of time sorting the data.

Slow queries for Parquet tables with `convert_legacy_hive_parquet_utc_timestamps=true`

The configuration setting `convert_legacy_hive_parquet_utc_timestamps=true` uses an underlying function that can be a bottleneck on high volume, highly concurrent queries due to the use of a global lock while loading time zone information. This bottleneck can cause slowness when querying Parquet tables, up to 30x for scan-heavy queries. The amount of slowdown depends on factors such as the number of cores and number of threads involved in the query.



Note:

The slowdown only occurs when accessing `TIMESTAMP` columns within Parquet files that were generated by Hive, and therefore require the on-the-fly timezone conversion processing.

Bug: [IMPALA-3316](#)

Severity: High

Workaround: Store the `TIMESTAMP` values as strings in one of the following formats:

- `yyyy-MM-dd`
- `yyyy-MM-dd HH:mm:ss`
- `yyyy-MM-dd HH:mm:ss.SSSSSSSS`

The date can have the 1-9 digits in the fractional part.

Impala implicitly converts such string values to `TIMESTAMP` in calls to date/time functions.

Slow DDL statements for tables with large number of partitions

DDL statements for tables with a large number of partitions might be slow.

Bug: [IMPALA-1480](#)

Workaround: Run the DDL statement in Hive if the slowness is an issue.

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0.

Interaction of File Handle Cache with HDFS Appends and Short-Circuit Reads

If a data file used by Impala is being continuously appended or overwritten in place by an HDFS mechanism, such as `hdfs dfs -appendToFile`, interaction with the file handle caching feature in CDH 5.13 / Impala 2.10 and higher could cause short-circuit reads to sometimes be disabled on some DataNodes. When a mismatch is detected between the cached file handle and a data block that was rewritten because of an append, short-circuit reads are turned off on the affected host for a 10-minute period.

The possibility of encountering such an issue is the reason why the file handle caching feature is currently turned off by default. See [Scalability Considerations for NameNode Traffic with File Handle Caching](#) for information about this feature and how to enable it.

Bug: [HDFS-12528](#)

Severity: High

Workaround: Verify whether your ETL process is susceptible to this issue before enabling the file handle caching feature. You can set the `impalad` configuration option `unused_file_handle_timeout_sec` to a time period that is shorter than the HDFS setting `dfs.client.read.shortcircuit.streams.cache.expiry.ms`. (Keep in mind that the HDFS setting is in milliseconds while the Impala setting is in seconds.)

Resolution: Fixed in HDFS 2.10 and higher. Use the new HDFS parameter `dfs.domain.socket.disable.interval.seconds` to specify the amount of time that short circuit reads are disabled on encountering an error. The default value is 10 minutes (600 seconds). It is recommended that you set `dfs.domain.socket.disable.interval.seconds` to a small value, such as 1 second, when using the file handle cache. Setting `dfs.domain.socket.disable.interval.seconds` to 0 is not recommended as a non-zero interval protects the system if there is a persistent problem with short circuit reads.

Impala Known Issues: Usability

These issues affect the convenience of interacting directly with Impala, typically through the Impala shell or Hue.

[Impala shell tarball is not usable on systems with setuptools versions where '0.7' is a substring of the full version string](#)

For example, this issue could occur on a system using setuptools version 20.7.0.

Bug: [IMPALA-4570](#)

Severity: High

Resolution: Fixed in CDH 5.10 / Impala 2.8 and higher.

Workaround: Change to a setuptools version that does not have 0.7 as a substring.

Unexpected privileges in show output

Due to a timing condition in updating cached policy data from Sentry, the `SHOW` statements for Sentry roles could sometimes display out-of-date role settings. Because Impala rechecks authorization for each SQL statement, this discrepancy does not represent a security issue for other statements.

Bug: [IMPALA-3133](#)

Severity: High

Resolution: Fixed in CDH 5.8.0 / Impala 2.6.0 and CDH 5.7.1 / Impala 2.5.1.

Less than 100% progress on completed simple SELECT queries

Simple `SELECT` queries show less than 100% progress even though they are already completed.

Bug: [IMPALA-1776](#)

Unexpected column overflow behavior with INT datatypes

Impala does not return column overflows as `NULL`, so that customers can distinguish between `NULL` data and overflow conditions similar to how they do so with traditional database systems. Impala returns the largest or smallest value in the range for the type. For example, valid values for a `tinyint` range from -128 to 127. In Impala, a `tinyint` with a value of -200 returns -128 rather than `NULL`. A `tinyint` with a value of 200 returns 127.

Bug: [IMPALA-3123](#)

Resolution: Fixed in CDH 5.8.0 and higher / Impala 2.6.0

Impala Known Issues: JDBC and ODBC Drivers

These issues affect applications that use the JDBC or ODBC APIs, such as business intelligence tools or custom-written applications in languages such as Java or C++.

ImpalaODBC: Can not get the value in the `SQLGetData(m-x th column)` after the `SQLBindCol(m th column)`

If the ODBC `SQLGetData` is called on a series of columns, the function calls must follow the same order as the columns. For example, if data is fetched from column 2 then column 1, the `SQLGetData` call for column 1 returns `NULL`.

Bug: [IMPALA-1792](#)

Workaround: Fetch columns in the same order they are defined in the table.

Impala Known Issues: Security

These issues are related to security features, such as Kerberos authentication, Sentry authorization, encryption, auditing, and redaction.

Missing authorization in Apache Impala may allow data injection

A malicious user who is authenticated with Kerberos may have unauthorized access to internal services used by Impala to transfer intermediate data during query execution. If details of a running query (e.g. query ID, query plan) are available, a user can craft some RPC requests with custom software to inject data into a running query or end query execution prematurely, leading to wrong results of the query.

Products affected: Apache Impala

Releases affected: CDH 5.15.0, CDH 5.15.1

Users affected: Any users of Impala who have configured Kerberos security

Date/time of detection: Aug 21, 2018

Detected by: Cloudera, Inc.

Severity (Low/Medium/High): 4.5 "Medium";

CVSS: 3.0/AV:A/AC:H/PR:L/UI:R/S:U/C:N/I:H/A:N/E:P/RL:T/RC:C/IR:H/MAV:A/MAC:H/MPR:L/MUI:R

Impact: Data injection may lead to wrong results of queries.

CVE: CVE-2018-11785

Immediate action required: Upgrade to a version which contains the fix or as a workaround, disable KRPC by setting `--use_krpc=false` in the **Impala Command Line Argument Advanced Configuration Snippet (Safety Valve)**. The workaround will disable some improvements in stability and performance implemented in CDH 5.15.0 for highly concurrent workloads.

Addressed in release/refresh/patch: CDH 5.15.2, CDH 5.16.1 and higher

Impala does not support Heimdal Kerberos

Heimdal Kerberos is not supported in Impala.

Bug: [IMPALA-7072](#)

Affected Versions: All versions of Impala

Transient kerberos authentication error during table loading

A transient Kerberos error can cause a table to get into a bad state with an error: Failed to load metadata for table.

Bug: [IMPALA-4712](#)

Severity: High

Workaround: Resolve the Kerberos authentication problem and run `INVALIDATE METADATA` on the affected table.

In Impala with Sentry enabled, `REVOKE ALL ON SERVER` does not remove the privileges granted with the `GRANT` option

If you grant a role the `ALL` privilege at the `SERVER` scope with the `WITH GRANT OPTION` clause, you cannot revoke the privilege. Although the `SHOW GRANT ROLE` command will show that the privilege has been revoked immediately after you run the command, the `ALL` privilege will reappear when you run the `SHOW GRANT ROLE` command after Sentry refreshes.

Immediate Action Required: Once the privilege has been granted, the only way to remove it is to delete the role.

Affected Versions: CDH 6.0.0, CDH 6.0.1, CDH 5.15.0, CDH 5.15.1, CDH 5.14.x and all prior releases

Fixed Versions: CDH 5.16.1, CDH 5.15.2

Cloudera Issue: TSB-341

Malicious user can gain unauthorized access to Kudu table data via Impala

A malicious user with `ALTER` permissions on an Impala table can access any other Kudu table data by altering the table properties to make it “external” and then changing the underlying table mapping to point to other Kudu tables. This violates and works around the authorization requirement that creating a Kudu external table via Impala requires an `ALL` privilege at the server scope. This privilege requirement for `CREATE` commands is enforced to precisely avoid this scenario where a malicious user can change the underlying Kudu table mapping. The fix is to enforce the same privilege requirement for `ALTER` commands that would make existing non-external Kudu tables external.

Bug: [IMPALA-5638](#)

Severity: High

Workaround: A temporary workaround is to revoke `ALTER` permissions on Impala tables.

Resolution: Fixed in CDH 5.11.2 / Impala 2.10 and higher

Catalog server's kerberos ticket gets deleted after 'ticket_lifetime' on SLES11

On SLES11, after 'ticket_lifetime', the kerberos ticket gets deleted by the Java krb5 library.

Bug: [IMPALA-6726](#)

Severity: High

Workaround: On Impala 2.11.0, set `--use_kudu_kinit=false` in Impala startup flag.

On Impala 2.12.0, set `--use_kudu_kinit=false` and `--use_krpc=false` in Impala startup flags.

Kerberos tickets must be renewable

In a Kerberos environment, the `impalad` daemon might not start if Kerberos tickets are not renewable.

Workaround: Configure your KDC to allow tickets to be renewed, and configure `krb5.conf` to request renewable tickets.

Impala does not allow the use of insecure clusters with public IPs

Starting in CDH 5.15 / Impala 2.12, Impala, by default, will only allow unencrypted or unauthenticated connections from trusted subnets: `127.0.0.0/8`, `10.0.0.0/8`, `172.16.0.0/12`, `192.168.0.0/16`, `169.254.0.0/16`.

Unencrypted or unauthenticated connections from publicly routable IPs will be rejected.

The trusted subnets can be configured using the `--trusted_subnets` flag. Set it to `'0.0.0.0/0'` to allow unauthenticated connections from all remote IP addresses. However, if network access is not otherwise restricted by a firewall, malicious users may be able to gain unauthorized access.

Impala user not added to `/etc/passwd` when LDAP is enabled

When using Impala with LDAP enabled, a user may hit the following:

```
Not authorized: Client connection negotiation failed: client connection to 127.0.0.1:27000:
SASL(-1): generic failure: All-whitespace username.
```

The following sequence can lead to the `impala` user not being created in `/etc/passwd` on some machines on the cluster.

- Time 1: The `impala` user is not in LDAP. Impala was installed on machine 1, and the user `impala` is created in `/etc/passwd`.
- Time 2: The `impala` user is added to LDAP.
- Time 3: A new machine is added to the cluster. When adding Impala service to this new machine, adding the `impala` user will fail as it already exists in LDAP.

The consequence is that the `impala` user doesn't exist in `/etc/passwd` on the new machine, leading to the error above.

Workaround: Manually edit `/etc/passwd` to add the `impala` user

Bug: [IMPALA-7585](#)

Affected Versions: CDH 5.15

Fixed Version: CDH 5.15.2, CDH 6.1

Kerberos authentication fails with the reverse DNS lookup disabled

Kerberos authentication does not function correctly if `rdns = false` is configured in `krb5.conf`. If the flag `rdns = false`, when Impala tries to match principals, it will fail because Kerberos receives a SPN (Service Principal Name) with an IP address in it, but Impala expects a principal with a FQDN in it.

Bug: [IMPALA-7298](#)

Affected Versions: CDH 5.15

Workaround: Set the following flags in `krb5.conf`:

- `dns_canonicalize_hostname = true`
- `rdns = true`

Fixed Version: CDH 5.15.1

System-wide `auth_to_local` mapping not applied correctly to Kudu service account

Due to system `auth_to_local` mapping, the principal may be mapped to some local name.

When running with Kerberos enabled, you may hit the following error message where `<random-string>` is some random string which doesn't match the primary in the Kerberos principal.

```
WARNINGS: TransmitData() to X.X.X.X:27000 failed: Remote error: Not authorized:
{username='<random-string>', principal='impala/redacted'} is not allowed to access
DataStreamService
```

Bug: [KUDU-2198](#)

Affected Versions: CDH 5.15 and higher

Workaround: Start Impala with the `--use_system_auth_to_local=false` flag to ignore the system-wide `auth_to_local` mappings configured in `/etc/krb5.conf`.

Kudu client in Impala FE doesn't renew/reacquire Kerberos tickets

When using Kudu with Impala on a secure cluster, the Kerberos ticket used by Kudu can expire even though the ticket used by Impala is renewed. Impala is then unable to access Kudu until the `impalad` daemon is restarted.

Cloudera Bug: CDH-63934

Affected versions: CDH 5.13.0, CDH 5.13.1

Severity: High

Workaround: Set a very long Kerberos ticket lifetime, for example 1 year.

Resolution: Upgrade to CDH 5.13.2 or higher.

Authorization error for SHOW CREATE VIEW

The SHOW CREATE VIEW statement returns an authorization error when the view references built-in functions.

Bug: [IMPALA-7325](#)

Affected Versions: CDH 5.15.1 and lower

Workaround: Execute following statement from Impala as an admin user. *role* refers to the role of the user who needs to execute the SHOW CREATE VIEW statement.

```
GRANT SELECT ON DATABASE `_impala_builtins` TO role;
```

Resolution: Fixed in CDH 5.16.1 and higher.

Impala Known Issues: Resources

These issues involve memory or disk usage, including out-of-memory conditions, the spill-to-disk feature, and resource management features.

Handling large rows during upgrade to CDH 5.13 / Impala 2.10 or higher

After an upgrade to CDH 5.13 / Impala 2.10 or higher, users who process very large column values (long strings), or have increased the `--read_size` configuration setting from its default of 8 MB, might encounter capacity errors for some queries that previously worked.

Bug: [IMPALA-6028](#)

Severity: High

Resolution: After the upgrade, follow the instructions in [Handling Large Rows During Upgrade to CDH 5.13 / Impala 2.10 or Higher](#) to check if your queries are affected by these changes and to modify your configuration settings if so.

Configuration to prevent crashes caused by thread resource limits

Impala could encounter a serious error due to resource usage under very high concurrency. The error message is similar to:

```
F0629 08:20:02.956413 29088 llvm-codegen.cc:111] LLVM hit fatal error: Unable to allocate
section memory!
terminate called after throwing an instance of
'boost::exception_detail::clone_impl<boost::exception_detail::error_info_injector<boost::thread_resource_error>
>'
```

Bug: [IMPALA-5605](#)

Severity: High

Workaround:

In CDH 5.14 and lower versions, configure each host running an `impalad` daemon with the following settings:

```
echo 2000000 > /proc/sys/kernel/threads-max
echo 2000000 > /proc/sys/kernel/pid_max
echo 8000000 > /proc/sys/vm/max_map_count
```

In CDH 5.15 and higher versions, it is unlikely that you will hit the thread resource limit. Configure each host running an `impalad` daemon with the following setting:

```
echo 8000000 > /proc/sys/vm/max_map_count
```

To make the above settings durable, refer to your OS documentation. For example, on RHEL 6.x:

1. Add the following line to `/etc/sysctl.conf`:

```
vm.max_map_count=8000000
```

2. Run the following command:

```
sysctl -p
```

Memory usage when `compact_catalog_topic` flag enabled

The efficiency improvement from [IMPALA-4029](#) can cause an increase in size of the updates to Impala catalog metadata that are broadcast to the `impalad` daemons by the `statedored` daemon. The increase in catalog update topic size results in higher CPU and network utilization. By default, the increase in topic size is about 5-7%. If the `compact_catalog_topic` flag is used, the size increase is more substantial, with a topic size approximately twice as large as in previous versions.

Bug: [IMPALA-5500](#)

Severity: Medium

Workaround: Consider setting the `compact_catalog_topic` configuration setting to `false` until this issue is resolved.

Resolution: Fixed in CDH 5.12.1 and CDH 5.13.0 and higher / Impala 2.10

Kerberos initialization errors due to high memory usage

On a kerberized cluster with high memory utilization, `kinit` commands executed after every `'kerberos_reinit_interval'` may cause out-of-memory errors, because executing the command involves a fork of the Impala process. The error looks similar to the following:

```
Failed to obtain Kerberos ticket for principal: <varname>principal_details</varname>  
Failed to execute shell cmd: 'kinit -k -t <varname>keytab_details</varname>',  
error was: Error(12): Cannot allocate memory
```

Bug: [IMPALA-2294](#)

Severity: High

Workaround:

The following command changes the `vm.overcommit_memory` setting immediately on a running host. However, this setting is reset when the host is restarted.

```
echo 1 > /proc/sys/vm/overcommit_memory
```

To change the setting in a persistent way, add the following line to the `/etc/sysctl.conf` file:

```
vm.overcommit_memory=1
```

Then run `sysctl -p`. No reboot is needed.

Resolution: Fixed in CDH 5.14.0 and higher / Impala 2.11.0

DROP TABLE PURGE on S3A table may not delete externally written files

A `DROP TABLE PURGE` statement against an S3 table could leave the data files behind, if the table directory and the data files were created with a combination of `hadoop fs` and `aws s3` commands.

Bug: [IMPALA-3558](#)

Severity: High

Resolution: The underlying issue with the S3A connector depends on the resolution of [HADOOP-13230](#).

Impala catalogd heap issues when upgrading to 5.7

The default heap size for Impala `catalogd` has changed in CDH 5.7 / Impala 2.5 and higher:

- Before 5.7, by default `catalogd` was using the JVM's default heap size, which is the smaller of 1/4th of the physical memory or 32 GB.
- Starting with CDH 5.7.0, the default `catalogd` heap size is 4 GB.

For example, on a host with 128GB physical memory this will result in `catalogd` heap decreasing from 32GB to 4GB. This can result in out-of-memory errors in `catalogd` and leading to query failures.

Severity: High

Workaround: Increase the `catalogd` memory limit as follows.

For schemas with large numbers of tables, partitions, and data files, the `catalogd` daemon might encounter an out-of-memory error. To increase the memory limit for the `catalogd` daemon:

1. Check current memory usage for the `catalogd` daemon by running the following commands on the host where that daemon runs on your cluster:

```

jcmd catalogd_pid VM.flags
jmap -heap catalogd_pid

```

2. Decide on a large enough value for the `catalogd` heap.

- On systems managed by Cloudera Manager, include this value in the configuration field **Java Heap Size of Catalog Server in Bytes** (Cloudera Manager 5.7 and higher), or **Impala Catalog Server Environment Advanced Configuration Snippet (Safety Valve)** (prior to Cloudera Manager 5.7). Then restart the Impala service.
- On systems not managed by Cloudera Manager, put the `JAVA_TOOL_OPTIONS` environment variable setting into the startup script for the `catalogd` daemon, then restart the `catalogd` daemon.

For example, the following environment variable setting specifies the maximum heap size of 8 GB.

```
JAVA_TOOL_OPTIONS="-Xmx8g"
```

3. Use the same `jcmd` and `jmap` commands as earlier to verify that the new settings are in effect.

Breakpad minidumps can be very large when the thread count is high

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Bug: [IMPALA-3509](#)

Severity: High

Workaround: Add `--minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the “hinted” size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

Parquet scanner memory increase after IMPALA-2736

The initial release of CDH 5.8 / Impala 2.6 sometimes has a higher peak memory usage than in previous releases while reading Parquet files.

CDH 5.8 / Impala 2.6 addresses the issue IMPALA-2736, which improves the efficiency of Parquet scans by up to 2x. The faster scans may result in a higher peak memory consumption compared to earlier versions of Impala due to the new column-wise row materialization strategy. You are likely to experience higher memory consumption in any of the following scenarios:

- Very wide rows due to projecting many columns in a scan.
- Very large rows due to big column values, for example, long strings or nested collections with many items.
- Producer/consumer speed imbalances, leading to more rows being buffered between a scan (producer) and downstream (consumer) plan nodes.

Bug: [IMPALA-3662](#)

Severity: High

Workaround: The following query options might help to reduce memory consumption in the Parquet scanner:

- Reduce the number of scanner threads, for example: `set num_scanner_threads=30`
- Reduce the batch size, for example: `set batch_size=512`
- Increase the memory limit, for example: `set mem_limit=64g`

Resolution: Fixed in CDH 5.10 / Impala 2.8.

Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the `impalad` daemon.

Bug: [IMPALA-691](#)

Workaround: To monitor overall memory usage, use the `top` command, or add the memory figures in the Impala web UI `/memz` tab to JVM memory usage shown on the `/metrics` tab.

Fix issues with the legacy join and agg nodes using `--enable_partitioned_hash_join=false` and `--enable_partitioned_aggregation=false`

Bug: [IMPALA-2375](#)

Workaround: Transition away from the “old-style” join and aggregation mechanism if practical.

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0.

Impala Known Issues: Correctness

These issues can cause incorrect or unexpected results from queries. They typically only arise in very specific circumstances.

DECIMAL AVG() can return incorrect results in Impala

When both an `AVG()` over a `DECIMAL` value and a `DISTINCT` aggregate (for example, `COUNT(DISTINCT ...)`) appear in the same `SELECT` statement, the `DECIMAL AVG()` can return incorrect results.

Bug: [IMPALA-5251](#)

Severity: High

Releases affected: CDH 5.11.0 only

Immediate action required:

- If you use Impala and are not yet on CDH 5.11.0, wait for CDH 5.11.1 to upgrade.
- If you use Impala and are in process of upgrading to CDH 5.11.0, or must upgrade before CDH 5.11.1 becomes available, contact Cloudera Support to get a patch to avoid this issue.
- If you are on CDH 5.11.0 and use Impala, contact Cloudera Support to get a patch to avoid this issue.
- If you do not use Impala, you are not affected.

Resolution: Fixed in CDH 5.12.0 and higher / Impala 2.9.0

Parquet scanner memory bug: I/O buffer is attached to output batch while scratch batch rows still reference it

Impala queries may return incorrect results when scanning plain-encoded string columns in uncompressed Parquet files. I/O buffers holding the string data are prematurely freed, leading to invalid memory reads and possibly non-deterministic results. This does not affect Parquet files that use a compression codec such as Snappy. Snappy is both strongly recommended generally and the default choice for Impala-written Parquet files.

How to determine whether a query might be affected:

- The query must reference `STRING` columns from a Parquet table.
- A selective filter on the Parquet table makes this issue more likely.
- Identify any uncompressed Parquet files processed by the query. Examine the `HDFS_SCAN_NODE` portion of a query profile that scans the suspected table. Use a query that performs a full table scan, and materializes the column values. (For example, `SELECT MIN(colname) FROM tablename.`) Look for “File Formats”. A value containing `PARQUET/NONE` means uncompressed Parquet.
- Identify any plain-encoded string columns in the associated table. Pay special attention to tables containing Parquet files generated through Hive, Spark, or other mechanisms outside of Impala, because Impala uses Snappy compression by default for Parquet files. Use `parquet-tools` to dump the file metadata. Note that a column could have several encodings within the same file (the column data is stored in several column chunks). Look for `VLE:PLAIN` in the output of `parquet-tools`, which means the values are plain encoded.

Bug: [IMPALA-4539](#)

Severity: High

Resolution: Fixed in CDH 5.10.0 / Impala 2.8 and higher

Workaround: Use Snappy or another compression codec for Parquet files.

ABS(n) where n is the lowest bound for the int types returns negative values

If the `abs()` function evaluates a number that is right at the lower bound for an integer data type, the positive result cannot be represented in the same type, and the result is returned as a negative number. For example, `abs(-128)` returns `-128` because the argument is interpreted as a `TINYINT` and the return value is also a `TINYINT`.

Bug: [IMPALA-4513](#)

Severity: High

Resolution: Fixed in CDH 5.14.0 and higher / Impala 2.11.0

Workaround: Cast the integer value to a larger type. For example, rewrite `abs(tinyint_col)` as `abs(cast(tinyint_col as smallint))`.

Java udf expression returning string in group by can give incorrect results.

If the `GROUP BY` clause included a call to a Java UDF that returned a string value, the UDF could return an incorrect result.

Bug: [IMPALA-4266](#)

Severity: High

Resolution: Fixed in CDH 5.10 / Impala 2.8 and higher.

Workaround: Rewrite the expression to concatenate the results of the Java UDF with an empty string call. For example, rewrite `my_hive_udf()` as `concat(my_hive_udf(), '')`.

Incorrect assignment of NULL checking predicate through an outer join of a nested collection.

A query could return wrong results (too many or too few NULL values) if it referenced an outer-joined nested collection and also contained a null-checking predicate (`IS NULL`, `IS NOT NULL`, or the `<=>` operator) in the `WHERE` clause.

Bug: [IMPALA-3084](#)

Severity: High

Resolution: Fixed in CDH 5.9.0 / Impala 2.7.0.

Incorrect result due to constant evaluation in query with outer join

An `OUTER JOIN` query could omit some expected result rows due to a constant such as `FALSE` in another join clause. For example:

```
explain SELECT 1 FROM alltypestiny a1
  INNER JOIN alltypesagg a2 ON a1.smallint_col = a2.year AND false
  RIGHT JOIN alltypes a3 ON a1.year = a1.bigint_col;
+-----+
| Explain String                                     |
+-----+
| Estimated Per-Host Requirements: Memory=1.00KB VCores=1 |
| 00:EMPTYSET                                       |
+-----+
```

Bug: [IMPALA-3094](#)

Severity: High

Workaround: None

Incorrect assignment of an inner join On-clause predicate through an outer join.

Impala might return incorrect results for queries that have the following properties:

- There is an `INNER JOIN` following a series of `OUTER JOINS`.
- The `INNER JOIN` has an `On-clause` with a predicate that references at least two tables that are on the nullable side of the preceding `OUTER JOINS`.

The following query demonstrates the issue:

```
select 1 from functional.alltypes a left outer join
  functional.alltypes b on a.id = b.id left outer join
  functional.alltypes c on b.id = c.id right outer join
  functional.alltypes d on c.id = d.id inner join functional.alltypes e
  on b.int_col = c.int_col;
```

The following listing shows the incorrect `EXPLAIN` plan:

```
+-----+
| Explain String                                     |
+-----+
| Estimated Per-Host Requirements: Memory=480.04MB VCores=4 |
| 14:EXCHANGE [UNPARTITIONED]                       |
| |                                                  |
| 08:NESTED LOOP JOIN [CROSS JOIN, BROADCAST]       |
+-----+
```

```

--13:EXCHANGE [BROADCAST]
|
| 04:SCAN HDFS [functional.alltypes e]
|   partitions=24/24 files=24 size=478.45KB
07:HASH JOIN [RIGHT OUTER JOIN, PARTITIONED]
|   hash predicates: c.id = d.id
|   runtime filters: RF000 <- d.id
--12:EXCHANGE [HASH(d.id)]
|
| 03:SCAN HDFS [functional.alltypes d]
|   partitions=24/24 files=24 size=478.45KB
06:HASH JOIN [LEFT OUTER JOIN, PARTITIONED]
|   hash predicates: b.id = c.id
|   other predicates: b.int_col = c.int_col    <--- incorrect placement; should be at
node 07 or 08
|   runtime filters: RF001 <- c.int_col
--11:EXCHANGE [HASH(c.id)]
|
| 02:SCAN HDFS [functional.alltypes c]
|   partitions=24/24 files=24 size=478.45KB
|   runtime filters: RF000 -> c.id
05:HASH JOIN [RIGHT OUTER JOIN, PARTITIONED]
|   hash predicates: b.id = a.id
|   runtime filters: RF002 <- a.id
--10:EXCHANGE [HASH(a.id)]
|
| 00:SCAN HDFS [functional.alltypes a]
|   partitions=24/24 files=24 size=478.45KB
09:EXCHANGE [HASH(b.id)]
|
| 01:SCAN HDFS [functional.alltypes b]
|   partitions=24/24 files=24 size=478.45KB
|   runtime filters: RF001 -> b.int_col, RF002 -> b.id
+-----+

```

Bug: [IMPALA-3126](#)

Severity: High

Resolution: Fixed in CDH 5.10.0 / Impala 2.8.0 and higher

Workaround: High

For some queries, this problem can be worked around by placing the problematic ON clause predicate in the WHERE clause instead, or changing the preceding OUTER JOINS to INNER JOINS (if the ON clause predicate would discard NULLS). For example, to fix the problematic query above:

```

select 1 from functional.alltypes a
  left outer join functional.alltypes b
    on a.id = b.id
  left outer join functional.alltypes c
    on b.id = c.id
  right outer join functional.alltypes d
    on c.id = d.id
  inner join functional.alltypes e
    where b.int_col = c.int_col

```

```

+-----+
| Explain String
+-----+

```

```

| Estimated Per-Host Requirements: Memory=480.04MB VCores=4
+-----+

```

```

14:EXCHANGE [UNPARTITIONED]
|
08:NESTED LOOP JOIN [CROSS JOIN, BROADCAST]
  --13:EXCHANGE [BROADCAST]
  |
  04:SCAN HDFS [functional.alltypes e]
    partitions=24/24 files=24 size=478.45KB
07:HASH JOIN [RIGHT OUTER JOIN, PARTITIONED]
  hash predicates: c.id = d.id
  other predicates: b.int_col = c.int_col
  runtime filters: RF000 <- d.id
  <-- correct assignment
  --12:EXCHANGE [HASH(d.id)]
  |
  03:SCAN HDFS [functional.alltypes d]
    partitions=24/24 files=24 size=478.45KB
06:HASH JOIN [LEFT OUTER JOIN, PARTITIONED]
  hash predicates: b.id = c.id
  --11:EXCHANGE [HASH(c.id)]
  |
  02:SCAN HDFS [functional.alltypes c]
    partitions=24/24 files=24 size=478.45KB
    runtime filters: RF000 -> c.id
05:HASH JOIN [RIGHT OUTER JOIN, PARTITIONED]
  hash predicates: b.id = a.id
  runtime filters: RF001 <- a.id
  --10:EXCHANGE [HASH(a.id)]
  |
  00:SCAN HDFS [functional.alltypes a]
    partitions=24/24 files=24 size=478.45KB
09:EXCHANGE [HASH(b.id)]
|
01:SCAN HDFS [functional.alltypes b]
  partitions=24/24 files=24 size=478.45KB
  runtime filters: RF001 -> b.id

```

Impala may use incorrect bit order with BIT_PACKED encoding

Parquet BIT_PACKED encoding as implemented by Impala is LSB first. The parquet standard says it is MSB first.

Bug: [IMPALA-3006](#)

Severity: High, but rare in practice because BIT_PACKED is infrequently used, is not written by Impala, and is deprecated in Parquet 2.0.

BST between 1972 and 1995

The calculation of start and end times for the BST (British Summer Time) time zone could be incorrect between 1972 and 1995. Between 1972 and 1995, BST began and ended at 02:00 GMT on the third Sunday in March (or second Sunday when Easter fell on the third) and fourth Sunday in October. For example, both function calls should return 13, but actually return 12, in a query such as:

```

select
  extract(from_utc_timestamp(cast('1970-01-01 12:00:00' as timestamp), 'Europe/London'),
    "hour") summer70start,
  extract(from_utc_timestamp(cast('1970-12-31 12:00:00' as timestamp), 'Europe/London'),
    "hour") summer70end;

```

Bug: [IMPALA-3082](#)

Severity: High

`parse_url()` returns incorrect result if @ character in URL

If a URL contains an @ character, the `parse_url()` function could return an incorrect value for the hostname field.

Bug: <https://issues.cloudera.org/browse/IMPALA-1170>IMPALA-1170

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0 and CDH 5.5.4 / Impala 2.3.4.

% escaping does not work correctly when occurs at the end in a LIKE clause

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

Bug: [IMPALA-2422](#)

ORDER BY rand() does not work.

Because the value for `rand()` is computed early in a query, using an ORDER BY expression involving a call to `rand()` does not actually randomize the results.

Bug: [IMPALA-397](#)

Resolution: Fixed in CDH 5.12.0 / Impala 2.9.0

Wrong results with correlated WHERE clause subquery inside a NULL-checking conditional function

Impala may generate an incorrect plan, and therefore incorrect results, for queries that have a correlated scalar subquery as a parameter to a NULL-checking conditional function such as `isnull()`.

Bug: [IMPALA-4373](#)

Severity: High

Workaround: None

Cannot execute IR UDF when single node execution is enabled

A UDF compiled into an LLVM IR bitcode module (`.bc`) would have undefined effects when native code generation was turned off, for example when Impala applied the single-node optimization for small queries.

Bug: [IMPALA-4432](#)

Severity: High

Resolution: In CDH 5.10 / Impala 2.8 and higher, Impala returns an error if the UDF cannot run because of this issue.

Workaround: Turn native code generation back on with the query option setting `DISABLE_CODEGEN=0`.

Duplicated column in inline view causes dropping null slots during scan

If the same column is queried twice within a view, NULL values for that column are omitted. For example, the result of `COUNT(*)` on the view could be less than expected.

Bug: [IMPALA-2643](#)

Workaround: Avoid selecting the same column twice within an inline view.

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0, CDH 5.5.2 / Impala 2.3.2, and CDH 5.4.10 / Impala 2.2.10.

Incorrect assignment of predicates through an outer join in an inline view.

A query involving an OUTER JOIN clause where one of the table references is an inline view might apply predicates from the ON clause incorrectly.

Bug: [IMPALA-1459](#)

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0, CDH 5.5.2 / Impala 2.3.2, and CDH 5.4.9 / Impala 2.2.9.

Crash: `impala::Coordinator::ValidateCollectionSlots`

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

Bug: [IMPALA-2603](#)

Incorrect assignment of On-clause predicate inside inline view with an outer join.

A query might return incorrect results due to wrong predicate assignment in the following scenario:

1. There is an inline view that contains an outer join
2. That inline view is joined with another table in the enclosing query block
3. That join has an On-clause containing a predicate that only references columns originating from the outer-joined tables inside the inline view

Bug: [IMPALA-2665](#)

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0, CDH 5.5.2 / Impala 2.3.2, and CDH 5.4.9 / Impala 2.2.9.

Wrong assignment of having clause predicate across outer join

In an `OUTER JOIN` query with a `HAVING` clause, the comparison from the `HAVING` clause might be applied at the wrong stage of query processing, leading to incorrect results.

Bug: [IMPALA-2144](#)

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0.

Wrong plan of NOT IN aggregate subquery when a constant is used in subquery predicate

A `NOT IN` operator with a subquery that calls an aggregate function, such as `NOT IN (SELECT SUM(...))`, could return incorrect results.

Bug: [IMPALA-2093](#)

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0 and CDH 5.5.4 / Impala 2.3.4.

Impala Known Issues: Metadata

These issues affect how Impala interacts with metadata. They cover areas such as the metastore database, the `COMPUTE STATS` statement, and the Impala `catalogd` daemon.

Concurrent catalog operations with heavy DDL workloads can cause queries with `SYNC_DDL` to fail fast

When Catalog Server is under a heavy load with concurrent catalog operations of long running DDLs, queries running with the `SYNC_DDL` query option can fail with the following message:

```
ERROR: CatalogException: Couldn't retrieve the catalog topic
version for the SYNC_DDL operation after 3 attempts.The operation has
been successfully executed but its effects may have not been
broadcast to all the coordinators.
```

The catalog operation is actually successful as the change has been committed to HMS and Catalog Server cache, but when Catalog Server notices a longer than expected time for it to broadcast the changes, it fails fast.

The coordinator daemons eventually sync up in the background.

Affected Versions: CDH versions from 5.15 to 6.1

Bug: [IMPALA-7961](#)

Resolution: Fixed in CDH 6.2.0 / Impala 3.2.

Catalogd may crash when loading metadata for tables with many partitions, many columns, and with incremental stats

Incremental stats use up about 400 bytes per partition for each column. For example, for a table with 20K partitions and 100 columns, the memory overhead from incremental statistics is about 800 MB. When serialized for transmission across the network, this metadata exceeds the 2 GB Java array size limit and leads to a `catalogd` crash.

Bugs: [IMPALA-2647](#), [IMPALA-2648](#), [IMPALA-2649](#)

Workaround: If feasible, compute full stats periodically and avoid computing incremental stats for that table. The scalability of incremental stats computation is a continuing work item.

Can't update stats manually via alter table after upgrading to CDH 5.2

Bug: [IMPALA-1420](#)

Workaround: On CDH 5.2, when adjusting table statistics manually by setting the `numRows`, you must also enable the Boolean property `STATS_GENERATED_VIA_STATS_TASK`. For example, use a statement like the following to set both properties with a single `ALTER TABLE` statement:

```
ALTER TABLE table_name SET TBLPROPERTIES('numRows'='new_value',
'STATS_GENERATED_VIA_STATS_TASK' = 'true');
```

Resolution: The underlying cause is the issue [HIVE-8648](#) that affects the metastore in Hive 0.13. The workaround is only needed until the fix for this issue is incorporated into a CDH release.

Impala Known Issues: Interoperability

These issues affect the ability to interchange data between Impala and other systems. They cover areas such as data types and file formats.

Queries Stuck on Failed HDFS Calls and not Timing out

In CDH 6.2 / Impala 3.2 and higher, if the following error appears multiple times in a short duration while running a query, it would mean that the connection between the `impalad` and the HDFS NameNode is in a bad state and hence the `impalad` would have to be restarted:

```
"hdfsOpenFile() for <filename> at backend <hostname:port> failed to finish before the
<hdfs_operation_timeout_sec> second timeout "
```

In CDH 6.1 / Impala 3.1 and lower, the same issue would cause Impala to wait for a long time or hang without showing the above error message.

Bug: [HADOOP-15720](#)

Affected Versions: All versions of Impala

Workaround: Restart the `impalad` in the bad state.

CREATE TABLE AS SELECT (CTAS) fails to write to HDFS

This issue can occur on clusters where HDFS NameNode high availability is enabled. The `CREATE TABLE AS SELECT` statement fails to open the HDFS file for writing. When this condition occurs an error is thrown in the Hue UI and logs that starts with "Failed to open HDFS file for writing:".

Severity: High

Workaround: To write data to HDFS on clusters where HDFS NameNode high availability is enabled, manually create and populate the table using the `CREATE TABLE` statement followed by `INSERT`.

DESCRIBE FORMATTED gives error on Avro table

This issue can occur either on old Avro tables (created prior to Hive 1.1 / CDH 5.4) or when changing the Avro schema file by adding or removing columns. Columns added to the schema file will not show up in the output of the `DESCRIBE FORMATTED` command. Removing columns from the schema file will trigger a `NullPointerException`.

As a workaround, you can use the output of `SHOW CREATE TABLE` to drop and recreate the table. This will populate the Hive metastore database with the correct column definitions.

**Warning:**

Only use this for external tables, or Impala will remove the data files. In case of an internal table, set it to external first:

```
ALTER TABLE table_name SET TBLPROPERTIES('EXTERNAL'='TRUE');
```

(The part in parentheses is case sensitive.) Make sure to pick the right choice between internal and external when recreating the table. See [Overview of Impala Tables](#) for the differences between internal and external tables.

Severity: High

Resolution: Fixed in CDH 5.8.2 and CDH 5.9.0 / Impala 2.6.2 and Impala 2.7.0

Avro Scanner fails to parse some schemas

The default value in Avro schema must match the first union type. For example, if the default value is `null`, then the first type in the `UNION` must be `"null"`.

Bug: [IMPALA-635](#)

Workaround: Swap the order of the fields in the schema specification. For example, use `["null", "string"]` instead of `["string", "null"]`. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Bug: [IMPALA-1024](#)

Severity: Remove trailing semicolon from the Avro schema.

Fix decompressor to allow parsing gzips with multiple streams

Currently, Impala can only read gzipped files containing a single stream. If a gzipped file contains multiple concatenated streams, the Impala query only processes the data from the first stream.

Bug: [IMPALA-2154](#)

Workaround: Use a different gzip tool to compress file to a single stream file.

Resolution: Fixed in CDH 5.7.0 / Impala 2.5.0.

Impala incorrectly handles text data when the new line character `\n\r` is split between different HDFS block

If a carriage return / newline pair of characters in a text table is split between HDFS data blocks, Impala incorrectly processes the row following the `\n\r` pair twice.

Bug: [IMPALA-1578](#)

Workaround: Use the Parquet format for large volumes of data where practical.

Resolution: Fixed in CDH 5.8.0 / Impala 2.6.0.

Invalid bool value not reported as a scanner error

In some cases, an invalid `BOOLEAN` value read from a table does not produce a warning message about the bad value. The result is still `NULL` as expected. Therefore, this is not a query correctness issue, but it could lead to overlooking the presence of invalid data.

Bug: [IMPALA-1862](#)

Resolution: Fixed in CDH 5.10.0 and higher / Impala 2.8.0

Incorrect results with basic predicate on CHAR typed column.

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Bug: [IMPALA-1652](#)

Workaround: Use the `RPAD()` function to blank-pad literals compared with CHAR columns to the expected length.

Impala Known Issues: Limitations

These issues are current limitations of Impala that require evaluation as you plan how to integrate Impala into your data management workflow.

Set limits on size of expression trees

Very deeply nested expressions within queries can exceed internal Impala limits, leading to excessive memory usage.

Bug: [IMPALA-4551](#)

Severity: High

Workaround: Avoid queries with extremely large expression trees. Setting the query option `disable_codegen=true` may reduce the impact, at a cost of longer query runtime.

Hue and BDR require separate parameters for Impala Load Balancer

Cloudera Manager supports a single parameter for specifying the Impala Daemon Load Balancer. However, because BDR and Hue need to use different ports when connecting to the load balancer, it is not possible to configure the load balancer value so that BDR and Hue will work correctly in the same cluster.

Workaround: To configure BDR with Impala, use the load balancer configuration either without a port specification or with the Beeswax port.

To configure Hue, use the **Hue Server Advanced Configuration Snippet (Safety Valve) for `impalad_flags`** to specify the load balancer address with the HiveServer2 port.

Affected Versions: CDH versions from 5.11 to 6.0

Bug: [OPSAPS-46641](#)

Impala Known Issues: Miscellaneous / Older Issues

These issues do not fall into one of the above categories or have not been categorized yet.

A failed CTAS does not drop the table if the insert fails.

If a `CREATE TABLE AS SELECT` operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Bug: [IMPALA-2005](#)

Workaround: Drop the new table manually after a failed `CREATE TABLE AS SELECT`.

Casting scenarios with invalid/inconsistent results

Using a `CAST()` function to convert large literal values to smaller types, or to convert special values such as NaN or INF, produces values not consistent with other database systems. This could lead to unexpected results from queries.

Bug: [IMPALA-1821](#)

Support individual memory allocations larger than 1 GB

The largest single block of memory that Impala can allocate during a query is 1 GiB. Therefore, a query could fail or Impala could crash if a compressed text file resulted in more than 1 GiB of data in uncompressed form, or if a string function such as `group_concat()` returned a value greater than 1 GiB.

Bug: [IMPALA-1619](#)

Resolution: Fixed in CDH 5.9.0 / Impala 2.7.0 and CDH 5.8.3 / Impala 2.6.3.

Impala Parser issue when using fully qualified table names that start with a number.

A fully qualified table name starting with a number could cause a parsing error. In a name such as `db.571_market`, the decimal point followed by digits is interpreted as a floating-point number.

Bug: [IMPALA-941](#)

Workaround: Surround each part of the fully qualified name with backticks (```).

Impala should tolerate bad locale settings

If the `LC_*` environment variables specify an unsupported locale, Impala does not start.

Bug: [IMPALA-532](#)

Workaround: Add `LC_ALL="C"` to the environment settings for both the Impala daemon and the Statestore daemon. See [Modifying Impala Startup Options](#) for details about modifying these environment settings.

Resolution: Fixing this issue would require an upgrade to Boost 1.47 in the Impala distribution.

Log Level 3 Not Recommended for Impala

The extensive logging produced by log level 3 can cause serious performance overhead and capacity issues.

Workaround: Reduce the log level to its default value of 1, that is, `GLOG_v=1`. See [Setting Logging Levels](#) for details about the effects of setting different logging levels.

Cloudera Distribution of Apache Kafka Known Issues

Flume Kafka Client Incompatible Changes in CDH 5.8

After an upgrade to CDH 5.8, Flume throws an error on startup and stops processing Kafka messages. This is caused by an incompatible configuration change: upgrade of the Kafka client from 0.8.2 to 0.9.0 in CDH 5.8.

Additionally, Flume may skip some Kafka messages because Kafka offset storage moves from ZooKeeper to Kafka with the client change. Without knowing about the previous offsets stored in ZooKeeper, the Flume Kafka client assumes the most recent offset, skipping over the existing messages that have not yet been read.

Products affected: Flume

Releases affected: CDH 5.8.0, CDH 5.8.1

Users affected: Users who have Flume configured with a Kafka source or channel. Flume Kafka sinks are not affected.

Severity (Low/Medium/High): High

Impact: Potential loss of messages consumed by the Flume agent from Kafka source, or loss of events in the Flume Kafka channel after upgrade

Immediate action required: If you are upgrading from a version lower than CDH 5.7.0, upgrade to CDH 5.7.x first. The Flume Kafka client version in lower versions of CDH does not support the migration properties required in the Flume configuration.

If you are currently running CDH 5.7.x, upgrade to CDH 5.8.2 or CDH 5.9.0 to resolve this issue. If you are upgrading to CDH 5.8.0 or CDH 5.8.1, see the following documentation to perform the steps required for Flume before upgrading: http://www.cloudera.com/documentation/enterprise/latest/topics/cdh_58_upgrade_kafka_flume_note.html

Addressed in release/refresh/patch: CDH 5.8.2 and higher, CDH 5.9.0 and higher

Cloudera Bug: TSB-173

Workaround: See [Upgrading to CDH 5.8 When Using the Flume Kafka Client](#)

Cloudera Manager 5.12.1 Installs the Latest Version of CDK powered By Apache Kafka by Default

Installing Kafka through Cloudera Manager 5.12.1 installs the latest available version of CDK Powered By Apache Kafka by default.

Workaround: In Cloudera Manager, go to **Hosts > Parcels > Configuration** and change the appropriate remote parcel repository to a parcel repository that hosts the required version of CDK Powered by Apache Kafka. Once the parcel repository is set, install the correct version of CDK Powered by Apache Kafka from the parcel. For more information regarding parcel configuration and installation, see [Parcel Configuration Settings](#), and [Installing or Upgrading Kafka from a Parcel](#).

Affected Versions: CDH 5.12.1

Fixed Versions: 5.13.0

Cloudera Issue: OPSAPS-43222

Apache Kudu Known Issues

The following are known bugs and issues in Kudu. Note that this list is not exhaustive, and is meant to communicate only the most important known issues.

Kudu Masters unable to join back after a restart

In a multi master Kudu environment, if a master is restarted or goes offline for a few minutes, it can occasionally have trouble joining the cluster on startup. For example, if this happens in case of three kudu masters, and one of the other two masters is stopped or dies during this time, then the overall Kudu cluster is down because the majority of the masters are not running.

This issue is resolved by the [KUDU-2748](#) upstream JIRA.

Products affected: Apache Kudu

Affected version:

- CDH 5.14.0, 5.14.2, 5.14.4
- CDH 5.15.0, 5.15.1, 5.15.2
- CDH 5.16.1, 5.16.2
- CDH 6.0.0, 6.0.1
- CDH 6.1.0, 6.1.1
- CDH 6.2.0, 6.2.1

Fixed version:

- CDH 6.3.0

For the latest update on this issue see the corresponding Knowledge article: [TSB 2020-442: Kudu Masters unable to join back after a restart](#)

Inconsistent rows returned from queries in Kudu

Due to KUDU-2463, upon restarting Kudu, inconsistent rows may be returned from tables that have not recently been written to, resulting in any of the following:

- multiple rows for the same key being returned
- deleted data being returned
- inconsistent results consistently being returned for the same query

If this happens, you have two options to resolve the conflicts: write to the affected Kudu partitions by:

- re-deleting the known and deleted data
- upserting the most up-to-date version of affected rows.

Products affected: Apache Kudu

Releases affected:

- CDH 5.12.2, 5.13.3, 5.14.4, 5.15.1, 5.16.1
- CDH 6.0.1, 6.1.0, 6.1.1

For the latest update on this issue see the corresponding Knowledge article: [TSB 2019-353: Inconsistent rows returned from queries in Kudu](#)

CFile Checksum Failure Causes Query to Fail

When a CFile checksum fails, for example, due to a underlying disk corruption, queries against the replica will fail with an error message, such as this:

```
Unable to advance iterator: Corruption: checksum error on CFile block
```

Workaround: Remove the corrupted replica from the tablet's Raft configuration. See [Kudu Troubleshooting Guide](#) for the detailed steps.

Affected Versions: CDH 6.0.x and lower

Apache Issue: [KUDU-2469](#)

The Kudu CLI tool crashes before commencing cluster rebalancing

The Kudu CLI crashes when running the 'kudu cluster rebalance' sub-command on the following platforms:

- RHEL/CentOS 7
- Ubuntu14.04 LTS (Trusty)
- SLES12

Affected Versions: CDH 5.16.0, CDH 5.16.1

Workaround: Specify an additional run-time flag `--move_single_replicas=disabled` when running the `kudu cluster rebalance` command. Specifying that flag explicitly disables balancing of non-replicated tables.

NOTE: As a good practice, non-replicated tables should not be used in production Kudu clusters.

Addressed in releases: CDH 5.16.2

Crash due to CHECK failure during compactions: `pv_delete_redo != nullptr`

Due to [KUDU-2233](#), tablet servers might crash on restart with a fatal log message like the following:

```
F1201 14:55:37.052140 10508 compaction.cc:756] Check failed: pv_delete_redo != nullptr
*
**
*** Check failure stack trace: ***
Wrote minidump to
/var/log/kudu/minidumps/kudu-tserver/215cde39-7795-0885-0b51038d-771d875e.dmp
*** Aborted at 1512161737 (unix time) try "date -d @1512161737" if you are using GNU
date ***
PC: @ 0x3ec3632625 (unknown)
*** SIGABRT (@0x3b98eec0000028e3) received by PID 10467 (TID 0x7f8b02c58700) from PID
10467; stack trace: ***
@ 0x3ec3a0f7e0 (unknown)
@ 0x3ec3632625 (unknown)
@ 0x3ec3633e05 (unknown)
0x1b53f59 (unknown)
@ 0x8b9f6d google::LogMessage::Fail()
@ 0x8bbe2d google::LogMessage::SendToLog()
@ 0x8b9aa9 google::LogMessage::Flush()
@ 0x8bc8cf google::LogMessageFatal::~LogMessageFatal()
@ 0x9db0fe kudu::tablet::FlushCompactionInput()
@ 0x9a056a kudu::tablet::Tablet::DoMergeCompactionOrFlush()
@ 0x9a372d kudu::tablet::Tablet::Compact()
@ 0x9bd8d1 kudu::tablet::CompactRowSetsOp::Perform()
@ 0x1b4145f kudu::MaintenanceManager::LaunchOp()
@ 0x1b8da06 kudu::ThreadPool::DispatchThread()
@ 0x1b888ea kudu::Thread::SuperviseThread()
@ 0x3ec3a07aa1 (unknown)
@ 0x3ec36e893d (unknown)
@ 0x0 (unknown)}}
```

After this crash occurs, the affected tablet server requires a patch in order to start again.

This crash is much more likely to occur on workloads involving deleting and re-inserting the same rows repeatedly.

Although data loss has not been observed as a result of this issue, data loss or corruption is possible. The integrity of a Kudu cluster's data can be verified using a `ksck` checksum scan:

```
sudo -u kudu kudu cluster ksck -checksum_scan <master addresses>
```

Products affected: Kudu

Releases affected:

- CDH 5.12.0, 5.12.1, 5.12.2
- CDH 5.13.0, CDH 5.13.1, CDH 5.13.2
- CDH 5.14.0, 5.14.1

Users affected: All users running Kudu, but the risk is much greater for workloads that involve deleting and re-inserting rows repeatedly.

Severity: High

Impact: The affected tablet server will not be able to start until a patch is applied. Data loss or corruption is possible.

Immediate action required: Upgrade to a version containing a fix.

If this is not possible, raising the minimum number of log segments retained to 2 will make the problem very unlikely to occur. To do this, go to the configuration page for Kudu and add `--log_min_segments_to_retain=2` to the Kudu Service Advanced Configuration Snippet (Safety Valve) for `gflagfile` option.

If the issue still occurs, it's necessary to patch the affected tablet server or upgrade to a version containing a fix.

Once upgraded to a version containing a fix, the `--log_min_segments_to_retain` setting should be removed.

Addressed in release: CDH5.13.3, CDH5.14.2

Partitions may be pruned incorrectly when range-partitioned on a primary key prefix

Without the fix from [KUDU-2173](#), the Kudu C++ client's partition pruner may incorrectly prune a partition when evaluating certain queries containing predicates on both the range partition columns and other columns that are part of the primary key. This may cause queries to fail to return some matching rows.

Specifically, the pruner mistakenly treats a range partition which is a proper prefix of a primary key as an exclusive bound, when in fact it is an inclusive bound if the remaining primary key column constraints are greater than the minimum value.

For example, given a simple schema:

```
CREATE TABLE t (
  a INT NOT NULL,
  b INT NOT NULL,
  PRIMARY KEY (a,b)
) PARTITION BY RANGE(a) PARTITIONS (
  ...
)
```

With the following range partition:

```
RANGE (a) PARTITION VALUES >= 10
```

And the following query:

```
SELECT * FROM t WHERE a < 11 AND b < 11
```

The partition pruner would exclude the above partition, even though it contains rows like (10, 1). Instead, the partition pruner treats `>= 10` as `> 10`. This wouldn't be a problem if the range partition was on column (b) instead, since that is not a prefix of the primary key. The same applies for (a, b), since that is the whole primary key and not a prefix.

Products affected: Kudu

Releases affected:

- CDH 5.10.0, 5.10.1, 5.10.2
- CDH 5.11.0, 5.11.1, 5.11.2
- CDH 5.12.0, 5.12.1
- CDH 5.13.0

Users affected: Direct Kudu users via its C++ or Python client. Indirect users via Impala.

Severity (Low/Medium/High): High

Impact: Certain queries won't return all the available matching rows and no error message will be displayed.

Immediate action required: Assess if any tables are currently affected by this bug.

1. Download the tool at <https://www.cloudera.com/content/dam/www/Support/Downloads/tsb-257.tar.gz>.
2. Extract the tool:

```
tar xzvf tsb-257.tar.gz
```

Run it:

```
java -jar tsb-257-check-0.1.0.jar <master-addresses> [table-name]
```

3. Capture the output. Depending on the output:

- If no tables are affected, ensure no new table is created with range partitioning on a proper prefix of the primary key until a new version is available for upgrade.
- If at least one table is affected, contact Cloudera Support to understand the next steps.

Addressed in release/refresh/patch:

- CDH 5.12.2
- CDH 5.13.1
- CDH 5.14.0
- CDH 5.15.0

C++ Client Fails to Re-acquire Authentication Token in Multi-master Clusters

A security-related issue can cause Impala queries to start failing on busy clusters in the following scenario:

- The cluster runs with the `--rpc_authentication` set as `optional` or `required`. The default is `optional`. Secure clusters use `required`.
- The cluster is using multiple masters.
- Impala queries happen frequently enough that the leader master connection to some `impalad` isn't idle-closed (more than 1 query per 65 seconds).
- The connection stays alive for longer than the authentication token timeout (1 week by default).
- A master leadership change occurs after the authentication token expiration.

Impala queries will start failing with errors in the `impalad` logs like:

```
I0904 13:53:08.748968 95857 client-internal.cc:283] Unable to determine the new leader
Master: Not authorized: Client connection negotiation failed: client connection to
10.164.44.13:7051: FATAL_INVALID_AUTHENTICATION_TOKEN: Not authorized: authentication
token expired
I0904 13:53:10.389009 95861 status.cc:125] Unable to open Kudu table: Timed out:
GetTableSchema timed out after deadline expired
@ 0x95b1e9 impala::Status::Status()
@ 0xff22d4 impala::KuduScanNodeBase::Open()
@ 0xff101e impala::KuduScanNode::Open()
@ 0xb73ced impala::FragmentInstanceState::Open()
```

```
@ 0xb7532b impala::FragmentInstanceState::Exec()
@ 0xb64ae8 impala::QueryState::ExecInstance()
@ 0xd15193 impala::Thread::SuperviseThread()
@ 0xd158d4 boost::detail::thread_data<>::run()
@ 0x129188a (unknown)
@ 0x7f717ceade25 start_thread
@ 0x7f717cbdb34d __clone
```

Impala shell queries will fail with a message like:

```
Unable to open Kudu table: Timed out: GetTableSchema timed out after deadline expired
```

Workaround:

- Restart the affected Impala Daemons. Restarting a daemon ensures the problem will not reoccur for at least the authentication token lifetime, which defaults to one week.
- Increase the authentication token lifetime (`--authn_token_validity_seconds`). Beware that raising this lifetime increases the window of vulnerability of the cluster if a client is compromised. It is recommended that you keep the token lifetime at one month maximum for a secure cluster. For unsecured clusters, a longer token lifetime is acceptable, and a 3 month lifetime is recommended.

Affected Versions: From CDH 5.11 through CDH 6.0

Fixed Versions: CDH 5.15.2, CDH 5.16.1

Apache Issue: [KUDU-2580](#)

Timeout Possible with Log Force Synchronization Option

If the Kudu master is configured with the `-log_force_fsync_all` option, tablet servers and clients will experience frequent timeouts, and the cluster may become unusable.

Longer Startup Times with a Large Number of Tablets

If a tablet server has a very large number of tablets, it may take several minutes to start up. It is recommended to limit the number of tablets per server to 1000 or fewer. The maximum allowed number of tablets is 2000 per server. Consider this limitation when pre-splitting your tables. If you notice slow start-up times, you can monitor the number of tablets per server in the web UI.

Kerberos Authentication Issues

- Kerberos authentication does not function correctly on hosts which contain capital letters in their hostname.
- Kerberos authentication does not function correctly if `rdns = false` is configured in `krb5.conf`.

Memory Issue with Fault Tolerant Scans

Unlike regular scans, fault tolerant scans will allocate all required memory when the scan begins rather than as it progresses. This can be significant for big tablets. Moreover, this memory usage isn't counted towards the tablet server's overall memory limit, raising the likelihood of the tablet server being out-of-memory killed by the kernel.

Affected Versions: CDH 6.2 / Kudu 1.9 and lower

Apache Issue: [KUDU-2466](#)

Confusing Descriptions for Kudu TLS/SSL Settings in Cloudera Manager

Descriptions in the Cloudera Manager Admin Console for TLS/SSL settings are confusing, and will be replaced in a future release. The correct information about the settings is in the Usage Notes column:

Field	Usage Notes
Kerberos Principal	Set to the default principal, <code>kudu</code> .

Field	Usage Notes
Enable Secure Authentication And Encryption	Select this checkbox to enable authentication <i>and</i> RPC encryption between all Kudu clients and servers, as well as between individual servers. Only enable this property after you have configured Kerberos.
Master TLS/SSL Server Private Key File (PEM Format)	Set to the path containing the Kudu master host's private key (PEM-format). This is used to enable TLS/SSL encryption (over HTTPS) for browser-based connections to the Kudu master web UI.
Tablet Server TLS/SSL Server Private Key File (PEM Format)	Set to the path containing the Kudu tablet server host's private key (PEM-format). This is used to enable TLS/SSL encryption (over HTTPS) for browser-based connections to Kudu tablet server web UIs.
Master TLS/SSL Server Certificate File (PEM Format)	Set to the path containing the signed certificate (PEM-format) for the Kudu master host's private key (set in Master TLS/SSL Server Private Key File). The certificate file can be created by concatenating all the appropriate root and intermediate certificates required to verify trust.
Tablet Server TLS/SSL Server Certificate File (PEM Format)	Set to the path containing the signed certificate (PEM-format) for the Kudu tablet server host's private key (set in Tablet Server TLS/SSL Server Private Key File). The certificate file can be created by concatenating all the appropriate root and intermediate certificates required to verify trust.
Master TLS/SSL Server CA Certificate (PEM Format)	Disregard this field.
Tablet Server TLS/SSL Server CA Certificate (PEM Format)	Disregard this field.
Enable TLS/SSL for Master Server	Enables HTTPS encryption on the Kudu master web UI.
Enable TLS/SSL for Tablet Server	Enables HTTPS encryption on the Kudu tablet server web UIs.

Apache Mahout Known Issues

The Scala module with Mahout does not support Java 8

Cloudera Bug: CDH-31784

Workaround: None

Apache Oozie Known Issues

[Apache Oozie Server Vulnerability CVE-2017-15712](#)

A vulnerability in the Oozie Server allows a cluster user to read private files owned by the user running the Oozie Server process.

Products affected: Oozie

Releases affected: All releases prior to CDH 5.12.0, CDH 5.12.0, CDH 5.12.1, CDH 5.12.2, CDH 5.13.0, CDH 5.13.1, CDH 5.14.0

Users affected: Users running the Oozie Server

Date/time of detection: November 13, 2017

Detected by: Daryn Sharp and Jason Lowe of Oath (formerly Yahoo! Inc)

Severity: (Low/Medium/High) High

Impact: The vulnerability allows a cluster user to read private files owned by the user running the Oozie Server process. The malicious user can construct a workflow XML file containing XML directives and configuration that reference sensitive files on the Oozie server host.

CVE: CVE-2017-15712

Immediate action required: Upgrade to release where the issue is fixed.

Addressed in release/refresh/patch: CDH 5.13.2 and higher, CDH 5.14.2 and higher, 5.15.0 and higher

Oozie does not consider timezone parameter when scheduling using cron-like syntax

Oozie workflows can be scheduled using coordinators. Coordinators can have start time, end time, and frequency parameters. All those have an effect on when and how often coordinator actions (occurrences) of the same coordinator job (definition) will schedule to run.

[There are three ways](#) to choose the frequency parameter:

1. Use a constant value (meaning number of minutes) like: `frequency="60"`. Note that because the number of minutes per day will vary in a daylight saving time zone, only values less than 1440 (minutes in a non-daylight-adjusting day) might be useful.
2. [Put an EL expression](#) like: `frequency=" ${coord:days(1)} "`. This will repeat the workflow once every day, based on the start time.
3. [Use a Cron expression](#) like: `frequency="10 9 * * *"`. This will schedule the coordinator application every day 9:10am.

If you create a scheduled job using either a constant value or an Expression Language expression while using the timezone parameter to supply a timezone with daylight saving time, Oozie schedules the coordinator actions considering daylight saving changes.

This is not the case when you use a Cron expression while using the timezone parameter to supply DST aware timezone. In this case, Oozie schedules the coordinator actions in a way that does not consider the timezone parameter. Only `oozie.processing.timezone` configuration value is considered configured as part of `oozie-site.xml`, and only for calculating the offset to GMT. `oozie-site.xml` affects the overall behavior for each coordinator job.

Affected Versions: CDH5.14.0 and lower.

Fixed in Versions: CDH5.16.1 and higher, CDH6.0.0 and higher

Bug: [OOZIE-2494](#)

Cloudera Bug: CDH-40279

Workaround: The workaround is use case specific. See the following Knowledge Base article and consider the listed caveats: [Oozie Coordinators Timezone and DST change handling caveats](#).

After enabling HA, Oozie may fail to start due to "NoSuchFieldError: EXTERNAL_PROPERTY

This issue happens in rare cases. Due to an incompatibility with the version of Jackson used by Oozie and Hive, and depending on the order that jars are loaded into Oozie's classpath, Oozie may fail to start.

Affected Versions: CDH5.9.3 and below.

Fixed in Versions: CDH 5.10.0 and higher.

Bug: [HIVE-1640](#)

Cloudera Bug: CDH-42408

Workaround: If using parcels:

1. Delete or move `/opt/cloudera/parcels/CDH/lib/oozie/libserver/hive-exec.jar` and `/opt/cloudera/parcels/CDH/lib/oozie/libtools/hive-exec.jar`.
2. Download `hive-exec-<cdh version>-core.jar` from the Cloudera repo and put it in `/opt/cloudera/parcels/CDH/lib/oozie/libserver/` and `/opt/cloudera/parcels/CDH/lib/oozie/libtools/`.
3. Download `kryo-2.22.jar` from the maven repo and put it in `/opt/cloudera/parcels/CDH/lib/oozie/libserver/` and `/opt/cloudera/parcels/CDH/lib/oozie/libtools/`.



Note: If using packages, use `/usr/lib/oozie` instead of `/opt/cloudera/parcels/CDH/lib/oozie/` in the above paths.

Oozie Web Console returns 500 error when Oozie server runs on JDK 8u75 or higher

The Oozie Web Console returns a 500 error when the Oozie server is running on JDK 8u75 and higher. The Oozie server still functions, and you can use the Oozie command line, REST API, Java API, or the Hue Oozie Dashboard to review status of those jobs.

Affected Versions: CDH5.x and higher, except for the releases listed below.

Fixed in Versions: CDH 5.5.5, 5.7.2, 5.8.2, 5.9.0 and above.

Bug: [OOZIE-2533](#)

Cloudera Bug: CDH-40362

Workaround: Use an earlier version of Java 8 or use the Hue Oozie Dashboard.

Oozie jobs fail (gracefully) on secure YARN clusters when JobHistory server is down

If the JobHistory server is down on a YARN (MRv2) cluster, Oozie attempts to submit a job, by default, three times. If the job fails, Oozie automatically puts the workflow in a `SUSPEND` state.

Affected Versions: CDH 5 Beta 1 and higher.

Cloudera Bug: CDH-14623

Workaround: When the JobHistory server is running again, use the `resume` command to tell Oozie to continue the workflow from the point at which it left off.

Oozie does not start when `oozie.email.smtp.auth` is disabled

If you enable SLA integration, and `oozie.email.smtp.auth` is disabled, Oozie throws a `NullPointerException` and fails to start.

Affected Versions: C5.5.1 and lower.

Bug: [OOZIE-2365](#)

Cloudera Bug: CDH-35331

Workaround: In Cloudera Manager, configure **Oozie Server Advanced Configuration Snippet (Safety Valve) for `oozie-site.xml`** as follows:

```
<property>
  <name>oozie.email.smtp.password</name>
  <value>none</value>
</property>
<property>
  <name>oozie.email.smtp.username</name>
  <value>none</value>
</property>
```

Oozie works with MapReduce or YARN, but not both

The Oozie server works with a MapReduce (MRv1) cluster or a YARN (MRv2) cluster, but not both at the same time.

Workaround: Use two different Oozie servers.

Apache Parquet Known Issues

Parquet file writes run out of memory if (number of partitions) times (block size) exceeds available memory

The Parquet output writer allocates one block for each table partition it is processing and writes partitions in parallel. The MapReduce or YARN task will run out of memory if (number of partitions) times (Parquet block size) is greater than the available memory.

Cloudera Bug: CDH-20157, CDH-20253

Workaround: None; if necessary, reduce the number of partitions in the table.

parquet-thrift cannot read Parquet data written by Hive

parquet-thrift cannot read Parquet data written by Hive, and parquet-avro will show an additional record level in lists named array_element.

Bug: [PARQUET-113](#)

Cloudera Bug: CDH-22189, CDH-22220

Workaround: None; arrays written by parquet-avro or parquet-thrift cannot currently be read by parquet-hive.

Apache Pig Known Issues

Missing results in Hive, Spark, Pig, Custom MapReduce jobs, and other Java applications when filtering Parquet data written by Impala

Apache Hive and Apache Spark rely on Apache Parquet's parquet-mr Java library to perform filtering of Parquet data stored in row groups. Those row groups contain statistics that make the filtering efficient without having to examine every value within the row group.

Recent versions of the parquet-mr library contain a bug described in [PARQUET-1217](#). This bug causes filtering to behave incorrectly if only some of the statistics for a row group are written. Starting in CDH 5.13, Apache Impala populates statistics in this way for Parquet files. As a result, Hive and Spark may incorrectly filter Parquet data that is written by Impala.

In CDH 5.13, Impala started writing Parquet's null_count metadata field without writing the min and max fields. This is valid, but it triggers the [PARQUET-1217](#) bug in the predicate push-down code of the Parquet Java library (parquet-mr). If the null_count field is set to a non-zero value, parquet-mr assumes that min and max are also set and reads them without checking whether they are actually there. If those fields are not set, parquet-mr reads their default value instead.

For integer SQL types, the default value is 0, so parquet-mr incorrectly assumes that the min and max values are both 0. This causes the problem when filtering data. Unless the value 0 itself matches the search condition, all row groups are discarded due to the incorrect min/max values, which leads to missing results.

Affected Products: The Parquet Java library (parquet-mr) and by extension, all Java applications reading Parquet files, including, but not limited to:

- Hive
- Spark
- Pig
- Custom MapReduce jobs

Affected Versions:

- CDH 5.13.0, 5.13.1, 5.13.2, and 5.14.0
- CDS 2.2 Release 2 Powered by Apache Spark and earlier releases on CDH 5.13.0 and later

Who Is Affected: Anyone writing Parquet files with Impala and reading them back with Hive, Spark, or other Java-based components that use the parquet-mr libraries for reading Parquet files.

Severity (Low/Medium/High): High

Impact: Parquet files containing null values for integer fields written by Impala produce missing results in Hive, Spark, and other Java applications when filtering by the integer field.

Immediate Action Required:

- **Upgrade**

You should upgrade to one of the fixed maintenance releases mentioned below.

- **Workaround**

This issue can be avoided at the price of performance by disabling predicate push-down optimizations:

- In Hive, use the following `SET` command:

```
SET hive.optimize.ppd = false;
```

- In Spark, disable the following configuration setting:

```
--conf spark.sql.parquet.filterPushdown=false
```

Addressed in the Following Releases:

- CDH 5.13.3 and higher
- CDH 5.14.2 and higher
- CDH 5.15.0 and higher
- CDS 2.3 Release 2 and higher

For the latest update on this issue, see the corresponding Knowledge Base article:

[TSB:2018-300: Missing results in Hive, Spark, Pig, and other Java applications when filtering Parquet data written by Impala](#)

Hive, Pig, and Sqoop 1 fail in MRv1 tarball installation because `/usr/bin/hbase` sets `HADOOP_MAPRED_HOME` to MRv2

This problem affects tarball installations only.

Cloudera Bug: CDH-6640

Resolution: Use workaround.

Workaround: If you are using MRv1, edit the following line in `/etc/default/hadoop` from

```
export HADOOP_MAPRED_HOME=/usr/lib/hadoop-mapreduce
```

to

```
export HADOOP_MAPRED_HOME=/usr/lib/hadoop-0.20-mapreduce
```

In addition, `/usr/lib/hadoop-mapreduce` must not exist in `HADOOP_CLASSPATH`.

Pig fails to read Parquet file (created with Hive) with a complex field if schema not specified explicitly

Affected Versions: CDH 5.1 - 5.12

Fixed in Versions: CDH 5.13 and higher

Cloudera Bug: CDH-15607

Workaround: Provide the schema of the fields in the `LOAD` statement.

Cloudera Search Known Issues

The current release includes the following known limitations:

Default Solr core names cannot be changed (limitation)

Although it is technically possible to give user-defined Solr core names during core creation, it is to be avoided in the context of Cloudera Search. Cloudera Manager expects core names in the default "collection_shardX_replicaY" format. Altering core names results in Cloudera Manager being unable to fetch Solr metrics for the given core and this, eventually, may corrupt data collection for co-located core, or even shard and server level charts.

Affected versions: All.

The `solrconfig.xml.secure` Template Does Not Enforce Apache Sentry Authorization

The `solrconfig.xml.secure` file generated by the `solrctl instancedir --generate` command does not enable Sentry authorization correctly. If you copied this file or renamed it to `solrconfig.xml` and used it for any collections, Sentry authorization is not being enforced on those collections.

Workaround: Modify the `solrconfig.xml` configuration file for each Sentry-protected collection as follows:

1. Download the collection configuration from ZooKeeper:

```
solrctl instancedir --get <config_name> /tmp/<config_name>
```

2. Edit the `/tmp/<config_name>/conf/solrconfig.xml` file as follows:

Change this line:

```
<updateRequestProcessorChain name="updateIndexAuthorization">
```

to this:

```
<updateRequestProcessorChain name="updateIndexAuthorization" default="true">
```

3. Upload the modified configuration to ZooKeeper:

```
solrctl instancedir --update <config_name> /tmp/<config_name>
```

Cloudera Bug: CDH-54101

Affected Versions: All CDH 5 versions except CDH 5.12.1 and higher

Fixed Versions: CDH 5.9.3, CDH 5.10.2, CDH 5.11.2, CDH 5.12.1

Collection Creation No Longer Supports Automatically Selecting A Configuration If Only One Exists

Before CDH 5.5.0, a collection could be created without specifying a configuration. If no `-c` value was specified, then:

- If there was only one configuration, that configuration was chosen.
- If the collection name matched a configuration name, that configuration was chosen.

Search for CDH 5.5.0 includes multiple built-in configurations. As a result, there is no longer a case in which only one configuration can be chosen by default.

To avoid this issue, explicitly specify the collection configuration to use by passing `-c configName` to `solrctl collection --create`.

Cloudera Bug: CDH-34050

Affected Versions: CDH 5.5.0 and higher

Creating an Instance Directory Fails If Matching Name Exists in `/solr/configs` Znode

The `solrctl instancedir --create` command checks for existing configurations under the `/solr/configs` ZooKeeper znode. If there is a znode with a name that matches the name of the config you are trying to create, the command fails.

Workaround: Use a different name for the instance directory, or make sure that the conflicting data in ZooKeeper is not needed and use `solrctl instancedir --update` instead.

Cloudera Bug: CDH-57036

Affected Versions: All CDH 5 versions except CDH 5.13.0 and higher

Fixed Versions: CDH 5.9.4, CDH 5.10.3, CDH 5.11.2, CDH 5.12.1, CDH 5.13.0

Configuration Templates Are Not Automatically Created on Existing SolrCloud Deployments

Configuration templates are only created when Solr is initialized. As a result, templates are not automatically available on existing SolrCloud deployments, even after upgrading to CDH 5.5.0.

If you do not need to retain information in your SolrCloud cluster, you can reinitialize solr using Cloudera Manager or using `solrctl init`. If you need to retain information in your cluster, there is no automated way to create the templates, but you can access the templates on a host with Solr installed at `/usr/lib/solr/templateName` for packages and `/opt/cloudera/parcels/CDH/lib/solr/templateName` for parcels. The templates can be uploaded to ZooKeeper using `instancedir` commands such as `solrctl instancedir --create templateName /path/to/templateName`, although they are not protected by ZooKeeper ACLs.

Cloudera Bug: CDH-34052

Affected Versions: CDH 5.5.0 and higher

Solr ZooKeeper ACLs Are Not Automatically Applied to Existing ZNodes

As of CDH 5.4, in Kerberos-enabled environments, ZooKeeper ACLs restrict access to Solr metadata stored in ZooKeeper to the solr user. This metadata cannot be modified by other users. These ACLs that limit access to the solr user are only applied automatically to new znodes.

This protection is not automatically applied to existing deployments.

To enable Solr ZooKeeper ACLs without retaining the existing cluster's Solr state, remove the solr znodes and reinitialize solr.

To remove solr znodes and reinitialize solr:

1. Using the `zookeeper-client`, enter the command `rmr /solr`.
2. Reinitialize Solr:
 - Select **Initialize Solr** in Cloudera Manager *OR*
 - Use `solrctl init`

To set ACLs using a script (CDH 5.7 and higher):

1. Create a `jaas.conf` file containing the following:

```
Client {
  com.sun.security.auth.module.Krb5LoginModule required
  useKeyTab=false
  useTicketCache=true
  principal="solr@EXAMPLE.COM";
};
```

Replace `EXAMPLE.COM` with your Kerberos realm name.

2. Set the `LOG4J_PROPS` environment variable to a `log4j.properties` file:

```
export LOG4J_PROPS=/etc/zookeeper/conf/log4j.properties
```

3. Set the `ZKCLI_JVM_FLAGS` environment variable:

```
export ZKCLI_JVM_FLAGS="-Djava.security.auth.login.config=/path/to/jaas.conf \  
-DzkACLProvider=org.apache.solr.common.cloud.ConfigAwareSaslZkACLProvider \  
-Droot.logger=INFO,console"
```

4. Authenticate as the `solr` user:

```
kinit solr@EXAMPLE.COM
```

Replace `EXAMPLE.COM` with your Kerberos realm name.

5. Run the `zkcli.sh` script as follows:

- **Cloudera Manager:**

```
/opt/cloudera/parcels/CDH/lib/solr/bin/zkcli.sh -zkhost zk01.example.com:2181 -cmd  
updateacls /solr
```

- **Unmanaged:**

```
/usr/lib/solr/bin/zkcli.sh -zkhost zk01.example.com:2181 -cmd updateacls /solr
```

Replace `zk01.example.com` with the hostname of a ZooKeeper server.

To enable Solr ZooKeeper ACLs while retaining the existing cluster's Solr state in versions lower than CDH 5.7, manually modify the existing `znode`'s ACL information. For example, using `zookeeper-client`, run the command `setAcl [path] sasl:solr:cdrwa,world:anyone:r`. This grants the `solr` user ownership of the specified path. Run this command for `/solr` and every `znode` under `/solr` except for the configuration `znodes` under and including `/solr/configs`.

Cloudera Bug: CDH-26416

Affected Versions: CDH 5.4.0 and higher

HBase Indexer ACLs Are Not Automatically Applied to Existing ZNodes

As of CDH 5.4, in Kerberos-enabled environments, ZooKeeper ACLs restrict access to Lily HBase Indexer metadata stored in ZooKeeper to `hbase` user. This metadata cannot be modified by other users. These ACLs that limit access to the `hbase` user are only applied automatically to new `znodes`.

This protection is not automatically applied to existing deployments.

To enable Lily HBase Indexer ACLs without retaining the existing cluster's Lily HBase Indexer state, turn off the Lily HBase Indexer, remove the `hbase-indexer` `znodes`, and then restart the Lily HBase Indexer.

To remove `hbase-indexer` `znodes` and reinitialize Lily HBase Indexer:

1. In Cloudera Manager, click



to the right of the Lily HBase Indexer service and select **Stop**.

2. Using the `zookeeper-client`, enter the command `rmr /ngdata`.

3. In Cloudera Manager, click



to the right of the Lily HBase Indexer service and select **Start**.

The Lily HBase Indexer automatically creates all required `znodes` when it is started.

To enable Lily HBase Indexer while retaining the existing HBase-Indexer state, manually modify the existing `znode`'s ACL information. For example, using `zookeeper-client`, run the command `setAcl`

[path]sas1:hbase:cdrwa,world:anyone:r. This grants the hbase user ownership of every znode under /ngdata (inclusive of /ngdata).



Note: This operation is not recursive, so creating a simple script may be helpful.

Cloudera Bug: CDH-26417

Affected Versions: CDH 5.4.0 and higher

[CrunchIndexerTool which includes Spark indexer requires specific input file format specifications](#)

If the `--input-file-format` option is specified with `CrunchIndexerTool`, then its argument must be `text`, `avro`, or `avroParquet`, rather than a fully qualified class name.

Cloudera Bug: CDH-22190

[Previously deleted empty shards may reappear after restarting the leader host](#)

It is possible to be in the process of deleting a collection when hosts are shut down. In such a case, when hosts are restarted, some shards from the deleted collection may still exist, but be empty.

Workaround: To delete these empty shards, manually delete the folder matching the shard. On the hosts on which the shards exist, remove folders under `/var/lib/solr/` that match the collection and shard. For example, if you had an empty shard 1 and empty shard 2 in a collection called `MyCollection`, you might delete all folders matching `/var/lib/solr/MyCollection{1,2}_replica*/`.

Cloudera Bug: CDH-20256

[The quickstart.sh file does not validate ZooKeeper and the NameNode on some operating systems](#)

The `quickstart.sh` file uses the `timeout` function to determine if ZooKeeper and the NameNode are available. To ensure this check can be complete as intended, the `quickstart.sh` determines if the operating system on which the script is running supports `timeout`. If the script detects that the operating system does not support `timeout`, the script continues without checking if the NameNode and ZooKeeper are available. If your environment is configured properly or you are using an operating system that supports `timeout`, this issue does not apply.

Workaround: This issue only occurs in some operating systems. If `timeout` is not available, a warning is displayed, but the `quickstart` continues and final validation is always done by the MapReduce jobs and Solr commands that are run by the `quickstart`.

Cloudera Bug: CDH-19923

[Field value class guessing and Automatic schema field addition are not supported with the MapReduceIndexerTool nor the HBaseMapReduceIndexerTool](#)

The `MapReduceIndexerTool` and the `HBaseMapReduceIndexerTool` can be used with a Managed Schema created via NRT indexing of documents or via the Solr Schema API. However, neither tool supports adding fields automatically to the schema during ingest.

Workaround: Define the schema before running the `MapReduceIndexerTool` or `HBaseMapReduceIndexerTool`. In non-schemaless mode, define in the schema using the `schema.xml` file. In schemaless mode, either define the schema using the Solr Schema API or index sample documents using NRT indexing before invoking the tools. In either case, Cloudera recommends that you verify that the schema is what you expect using the List Fields API command.

Cloudera Bug: CDH-26856

[The “Browse” and “Spell” Request Handlers are not enabled in schemaless mode](#)

The “Browse” and “Spell” Request Handlers require certain fields be present in the schema. Since those fields cannot be guaranteed to exist in a Schemaless setup, the “Browse” and “Spell” Request Handlers are not enabled by default.

Workaround: If you require the “Browse” and “Spell” Request Handlers, add them to the `solrconfig.xml` configuration file. Generate a non-schemaless configuration to see the usual settings and modify the required fields to fit your schema.

Cloudera Bug: CDH-19407

Enabling blockcache writing may result in unusable indexes

It is possible to create indexes with `solr.hdfs.blockcache.write.enabled` set to `true`. Such indexes may appear corrupt to readers, and reading these indexes may irreversibly corrupt indexes. Blockcache writing is disabled by default.

Workaround: Do not enable blockcache writing.

Cloudera Bug: CDH-17978

Solr fails to start when Trusted Realms are added for Solr into Cloudera Manager

Cloudera Manager generates name rules with spaces as a result of entries in the Trusted Realms, which do not work with Solr. This causes Solr to not start.

Workaround: Do not use the Trusted Realm field for Solr in Cloudera Manager. To write your own name rule mapping, add an environment variable `SOLR_AUTHENTICATION_KERBEROS_NAME_RULES` with the mapping. See [Cloudera Security](#) for more information.

Cloudera Bug: CDH-17289

Lily HBase batch indexer jobs fail to launch

A symptom of this issue is an exception similar to the following:

```
Exception in thread "main" java.lang.IllegalAccessError: class
com.google.protobuf.ZeroCopyLiteralByteString cannot access its superclass
com.google.protobuf.LiteralByteString
    at java.lang.ClassLoader.defineClass1(Native Method)
    at java.lang.ClassLoader.defineClass(ClassLoader.java:792)
    at java.security.SecureClassLoader.defineClass(SecureClassLoader.java:142)
    at java.net.URLClassLoader.defineClass(URLClassLoader.java:449)
    at java.net.URLClassLoader.access$100(URLClassLoader.java:71)
    at java.net.URLClassLoader$1.run(URLClassLoader.java:361)
    at java.net.URLClassLoader$1.run(URLClassLoader.java:355)
    at java.security.AccessController.doPrivileged(Native Method)
    at java.net.URLClassLoader.findClass(URLClassLoader.java:354)
    at java.lang.ClassLoader.loadClass(ClassLoader.java:424)
    at java.lang.ClassLoader.loadClass(ClassLoader.java:357)
    at org.apache.hadoop.hbase.protobuf.ProtobufUtil.toScan(ProtobufUtil.java:818)
    at
org.apache.hadoop.hbase.mapreduce.TableMapReduceUtil.convertScanToString(TableMapReduceUtil.java:433)
    at
org.apache.hadoop.hbase.mapreduce.TableMapReduceUtil.initTableMapperJob(TableMapReduceUtil.java:186)
    at
org.apache.hadoop.hbase.mapreduce.TableMapReduceUtil.initTableMapperJob(TableMapReduceUtil.java:147)
    at
org.apache.hadoop.hbase.mapreduce.TableMapReduceUtil.initTableMapperJob(TableMapReduceUtil.java:270)
    at
org.apache.hadoop.hbase.mapreduce.TableMapReduceUtil.initTableMapperJob(TableMapReduceUtil.java:100)
    at
com.ngdata.hbaseindexer.mr.HBaseMapReduceIndexerTool.run(HBaseMapReduceIndexerTool.java:124)
    at
com.ngdata.hbaseindexer.mr.HBaseMapReduceIndexerTool.run(HBaseMapReduceIndexerTool.java:64)
    at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:70)
```

```

at
com.ngdata.hbaseindexer.mr.HBaseMapReduceIndexerTool.main(HBaseMapReduceIndexerTool.java:51)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57)
at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
at java.lang.reflect.Method.invoke(Method.java:606)
at org.apache.hadoop.util.RunJar.main(RunJar.java:212)

```

This is because of an optimization introduced in [HBASE-9867](#) that inadvertently introduced a classloader dependency. In order to satisfy the new classloader requirements, `hbase-protocol.jar` must be included in Hadoop's classpath. This can be resolved on a per-job launch basis by including it in the `HADOOP_CLASSPATH` environment variable when you submit the job.

Workaround: Run the following command before issuing Lily HBase MapReduce jobs. Replace the `.jar` file names and filepaths as appropriate.

```

$ export HADOOP_CLASSPATH=</path/to/hbase-protocol>.jar; hadoop jar <MyJob>.jar
<MyJobMainClass>

```

Cloudera Bug: CDH-16539

[Users with insufficient Solr permissions may receive a "Page Loading" message from the Solr Web Admin UI](#)

Users who are not authorized to use the Solr Admin UI are not given page explaining that access is denied, and instead receive a web page that never finishes loading.

Workaround: None

Cloudera Bug: CDH-58276

[Using MapReduceIndexerTool or HBaseMapReduceIndexerTool multiple times may produce duplicate entries in a collection.](#)

Repeatedly running the `MapReduceIndexerTool` on the same set of input files can result in duplicate entries in the Solr collection. This occurs because the tool can only insert documents and cannot update or delete existing Solr documents. This issue does not apply to the `HBaseMapReduceIndexerTool` unless it is run with more than zero reducers.

Workaround: To avoid this issue, use `HBaseMapReduceIndexerTool` with zero reducers. This must be done without Kerberos.

Cloudera Bug: CDH-15441

[Deleting collections may fail if hosts are unavailable.](#)

It is possible to delete a collection when hosts that host some of the collection are unavailable. After such a deletion, if the previously unavailable hosts are brought back online, the deleted collection may be restored.

Workaround: Ensure all hosts are online before deleting collections.

Cloudera Bug: CDH-58694

[Saving search results is not supported.](#)

Cloudera Search does not support the ability to save search results.

Workaround: None

Cloudera Bug: CDH-21162

[HDFS Federation is not supported.](#)

Cloudera Search does not support HDFS Federation.

Workaround: None

Cloudera Bug: CDH-11357

Apache Sentry Known Issues

Sentry doesn't support Kafka topic name with more than 64 characters

A Kafka topic name can have 249 characters, but Sentry only supports topic names up to 64 characters.

Workaround: Keep Kafka topic names to 64 characters or less.

Affected Versions: All CDH 5.x and 6.x versions

Cloudera Issue: CDH-64317

When granting privileges, a single transaction per grant causes long delays

Sentry takes a long time to grant or revoke a large number of column-level privileges that are requested in a single statement. For example if you execute the following command:

```
GRANT SELECT(col1, col2, ...) ON TABLE table1;
```

Sentry applies the grants to each column separately and the refresh process causes long delays.

Workaround: Split the grant statement up into smaller chunks. This prevents the refresh process from causing delays.

Affected Versions:

- CDH: 5.14.4
- CDH: 5.15.1
- CDH: 5.16.0
- CDH: 6.1.0

Fixed Versions:

- CDH 5.16.1 and above
- CDH 6.2.0 and above

Cloudera Issue: CDH-74982

Snapshot doesn't complete with null SDS entries

If you have a partition with null SDS entries, a full snapshot will never complete.

Workaround: Delete the null SDS entries from HMS and restart Sentry.

Affected Versions: CDH 5.16.1

Fixed Versions: CDH 5.16.1 and above

Cloudera Issue: CDH-75158

SHOW ROLE GRANT GROUP raises exception for a group that was never granted a role

If you run the command SHOW ROLE GRANT GROUP for a group that has never been granted a role, beeline raises an exception. However, if you run the same command for a group that does not have any roles, but has at one time been granted a role, you do not get an exception, but instead get an empty list of roles granted to the group.

Workaround: Adding a role will prevent the exception.

Affected Versions: CDH 5.16.1

Cloudera Issue: CDH-71694

Sentry allows the ALTER TABLE EXCHANGE PARTITION operation on a restricted database

If a user has ALL permissions on a database, the [ALTER TABLE EXCHANGE PARTITION](#) command allows the user to move partitions from a table that the user does not have access to. For example, if a user has ALL permissions on database A, but no permissions on database B, the user can create a table with a schema in database A that is identical to a table in database B. The user can then move partitions from database B into database A, which allows the user to view restricted data and remove that data from the source database.

After you upgrade to a version of CDH listed in the "Addressed in release" section below, a user that tries to use the EXCHANGE PARTITION command to move a partition from a restricted database will receive a "No valid privileges" error.

Products affected: Hive services running Sentry

Affected Versions:

- CDH 5.13.x and below
- CDH 5.14.0, 5.14.2, 5.14.3
- CDH 5.15.0

Users affected: Hive users running Sentry

Date/time of detection: May 10, 2018

Severity (Low/Medium/High): 8.1 High ([CVSS:3.0/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:N](#))

Impact: Sensitive data exposure

CVE: CVE-2018-8028

Immediate action required: Upgrade to a version of CDH with the fix.

Addressed in release/refresh/patch:

- CDH 5.14.4
- CDH 5.15.1
- CDH 5.16.1
- CDH 6.0.0

For the latest update on this issue see the corresponding Knowledge article:

[TSB 2018-312: Sentry allows the "ALTER TABLE EXCHANGE PARTITION" operation on a restricted database](#)

Performance Issue: MPath is queried for each MAAuthzPathsMapping in full snapshot

MAAuthzPathsMapping contains list of MPath instances. From log message, when getting path full snapshot at SentryStore.retrieveFullPathsImageCore(), DataNucleus issues a query for all MPath instances associated with each MAAuthzPathsMapping. Therefore, getting full path image may take a very long time.

Workaround: The solution is to get MPath in a batch when getting full path image.

Affected Versions: CDH 5.13 and below

Apache Issue: [SENTRY-2184](#)

Cloudera Issue: None

The REVOKE GRANT OPTION also revokes privileges with no grant option

When you use the REVOKE GRANT OPTION command to revoke the GRANT privilege from a role, Hive revokes the privilege that has the GRANT OPTION and if the role also has the privilege without the grant option, that is revoked as well.

For example, if a role has SELECT and SELECT WITH GRANT OPTION privileges, and you run the REVOKE GRANT OPTION command, the role will have no SELECT privileges on that object. Both the SELECT and SELECT WITH GRANT OPTION privileges will be revoked.

Workaround: Revoke the GRANT OPTION and the privilege, and then grant the privilege.

Affected Versions: All CDH 5.x versions

Apache Issue: [SENTRY-2402](#)

Cloudera Issue: CDH-72979

After upgrade, HDFS ACLs are not synched due to Sentry taking too many snapshots

The schema change that was introduced in CDH 5.13.0 requires Sentry to take an HMS snapshot when you upgrade. Sentry might take too many snapshots, resulting in a period of time in which ACLs are not synched with HDFS. This can cause job failures.

Note that this issue is similar to [CDH-67922](#), but has a different root cause.

Who is Affected: Users that currently have any CDH version below 5.13.0 and are upgrading to any of the affected versions. For example, users upgrading from CDH 5.12.2 to CDH 5.13.0 are affected, but users upgrading from 5.12.2 to 5.13.3 are not affected.

Action Required: Cloudera recommends that if you have Sentry installed, you do not upgrade to any of the affected versions. Instead, upgrade to one of the fixed versions.

Affected Versions:

- CDH 5.13.0, 5.13.1, 5.13.2
- CDH 5.14.0, 5.14.1

Fixed Versions:

- CDH 5.13.3
- CDH 5.14.2, 5.14.3
- CDH 5.15.0 and above

Apache Issue: [SENTRY-2210](#)

Cloudera Issue: CDH-70480

After upgrade, HDFS ACLs are not synched due to Sentry taking too long to retrieve the snapshot from the database

The schema change that was introduced in 5.13.0 requires Sentry to take an HMS snapshot and save it to the Sentry database when you upgrade. When the NameNode service starts after the upgrade and asks for a path update, Sentry must retrieve a full snapshot from the Sentry database and send it to the NameNode. In the Sentry database, AUTHZ_PATH does not have an index on the foreign key AUTHZ_OBJ_ID. This can cause the snapshot retrieval from the Sentry database to take a long time, resulting in ACLs not being synched with HDFS until the snapshot is complete. After Sentry retrieves the full snapshot and sends it to the NameNode, Sentry only needs to retrieve and send delta changes.

Note that this issue is similar to [CDH-70480](#), but has a different root cause.

Who is Affected: Users that currently have any CDH version below 5.13.0 and are upgrading to any of the affected versions. For example, users upgrading from CDH 5.12.2 to CDH 5.13.3 are affected. MySQL users are not affected because MySQL automatically creates the index for the foreign key.

Action Required: After you upgrade, manually add the index to your schema with the appropriate command for the database:

- Oracle:

```
CREATE INDEX "AUTHZ_PATH_FK_IDX" ON "AUTHZ_PATH" ("AUTHZ_OBJ_ID");
```

- PostgreSQL:

```
CREATE INDEX "AUTHZ_PATH_FK_IDX" ON "AUTHZ_PATH" USING btree ("AUTHZ_OBJ_ID");
```

- Apache Derby:

```
CREATE INDEX AUTHZ_PATH_FK_IDX ON AUTHZ_PATH (AUTHZ_OBJ_ID);
```

- IBM Db2:

```
CREATE INDEX AUTHZ_PATH_FK_IDX ON AUTHZ_PATH (AUTHZ_OBJ_ID);
```

Affected Versions:

- CDH 5.13.0, 5.13.1, 5.13.2, 5.13.3
- CDH 5.14.0, 5.14.1, 5.14.2, 5.14.3
- CDH 5.15.0

Fixed Versions: CDH 5.16.1

Cloudera Issue: CDH-67922

You must use a relational database with Sentry HA

A flat-file cannot scale in a meaningful way with regards to search, replace, add, and delete operations - but a relational database can.

Workaround: Use a relational database.

Affected Versions: CDH 5.13.0 and above

Cloudera Issue: None

GRANT/REVOKE operations could fail if there are too many concurrent requests

Under a significant workload, Grant/Revoke operations can have issues.

Workaround: Given that grant and revoke actions are largely an administrative concern, plan en masse privilege changes accordingly.

Affected Versions: 5.13.0 and above

Apache Issue: [SENTRY-1855](#)

Cloudera Issue: CDH-56553

Creating large set of Sentry roles results in performance problems

Using more than a thousand roles/permissions may cause significant performance problems.

Workaround: In general it is better for users to have only a few roles associated as possible; then allowing for administrator roles to have very broad privileges.

Affected Versions: All CDH 5.x versions

Cloudera Issue: CDH-59010

Sentry takes a long time to finish HMS sync on larger object count (>2M)

When there are many Hive objects (tables and partitions) and many is in millions, it may take Sentry a few minutes (or a few tens of minutes) to fully synchronize table information between HMS and Sentry and during this time HDFS sync would not be functional. This delay is usually a one-time cost during the upgrade to 5.13 and should not be a problem during Sentry server restart.

Workaround: Be cognizant of the expected time to processes many objects. Scheduling the installation/upgrade during a low-usage time and when the system can simply process all of the objects would be best.

Affected Versions: 5.13.0 and above

Cloudera Issue: CDH-59310

Sentry process crashes with OOM upon reaching array length limit

To prevent Sentry from crashing with OOM upon reaching array length limit, use a Java version that has that has [JDK-8055949](#) fixed. This also affects HMS when Sentry is enabled and there are a lot of partitions.

Workaround: Use a Java version that has that has [JDK-8055949](#) fixed.

Affected Versions: All CDH 5 versions

Cloudera Issue: None

YARN Dynamic Resource Pools Do Not Work with Hive When Sentry Is Enabled

Hive jobs are not submitted into the correct YARN queue when Hive is using Sentry because Hive does not use the YARN API to resolve the user or group of the job's original submitter. This causes the job to be placed in a queue using the placement rules based on the Hive user. The HiveServer2 fair scheduler queue mapping used for "non-impersonation" mode does not handle the primary-secondary queue mappings correctly.

Products Affected: Hive, YARN, Sentry

Affected Versions:

This issue is fixed by [HIVE-8634](#), which is available in CDH 5.3 and higher. However, in CDH 5.3, 5.4, 5.5, 5.6, and 5.7, you must configure Hive to load the scheduler configuration into HiveServer2. In CDH 5.8 or higher managed clusters, Cloudera Manager automatically deploys the fair scheduler configuration to HiveServer2, but you must make sure that your configuration complies with the [Placement Rule](#) and [Submission Access](#) requirements described below.



Important: If you are using CDH 5.0.x, 5.1.0 to 5.1.4, or 5.2.0, you must upgrade to a later release that includes the fix ([HIVE-8634](#)).

Apache Issue: [HIVE-8634](#)

Resolution: Use the workaround for your version of CDH.

Workaround for CDH 5.3, 5.4, 5.5, 5.6, 5.7:

Also see [Placement Rule Requirement](#) and [Job Submission Access Requirement](#) for additional configuration information.

- **Managed clusters:**

1. In the Cloudera Manager Admin Console, click **Cluster *n***, select the **YARN** service, and click the **Instances** tab.
2. In the YARN Instances page, click the **ResourceManager(Active)** role, and then click the **Processes** tab.
3. In the ResourceManager Processes page, click **fair-scheduler.xml**, which opens in a new browser tab. Copy the contents of the **fair-scheduler.xml** page.
4. Create a new file named `fair-scheduler.xml` and paste the contents of the **fair-scheduler.xml** browser page into it. Save the file.
5. In a terminal window, create a new directory on the HiveServer2 host where you can store the `fair-scheduler.xml` file you created in the previous step. For example, `/etc/hive/fsxml`.



Important: Do not place the `fair_scheduler.xml` file into the standard Hive configuration directory. That directory is managed by Cloudera Manager and the file might be removed when you change other configuration settings.

6. Upload the `fair-scheduler.xml` file to the directory you created in the previous step.
7. In the Cloudera Manager Admin Console, click **Cloudera Manager** to return to the home page, and click **Hive > Configuration**.
8. On the Hive service Configuration page, select **Service-Wide** scope and **Advanced** category.
9. In the panel on the right, scroll down until you find the **Hive Service Advanced Configuration Snipped (Safety Valve) for hive-site.xml**, and add the `yarn.scheduler.fair.allocation.file` property with the path

to the `fair-scheduler.xml` file you created in the previous steps. For example, if you put the `fair-scheduler.xml` file in the `/etc/hive/fsxml` directory, add the following:

```
<property>
  <name>yarn.scheduler.fair.allocation.file</name>
  <value>/etc/hive/fsxml/fair-scheduler.xml</value>
</property>
```

10 Click **Save Changes**.

11 Restart the Hive service.



Note: An updated configuration file must be deployed each time a change is made in the fair scheduler configuration. The new file is read automatically by HiveServer2.

- **Unmanaged clusters:**

1. Log in to a node that runs the ResourceManager.
2. Navigate to the `$HADOOP_HOME/conf/` directory and locate the `fair-scheduler.xml` file.
3. Copy the `fair-scheduler.xml` file to the HiveServer2 node in the `/etc/hive/conf` directory.
4. Open the `/etc/hive/conf/hive-site.xml` file and add the following property:

```
<property>
  <name>yarn.scheduler.fair.allocation.file</name>
  <value>/etc/hive/conf/fair-scheduler.xml</value>
</property>
```

5. Save the `hive-site.xml` file.

6. Restart the Hive service:

```
$ sudo service hive-server2 stop
$ sudo service hive-server2 start
```



Note: An updated configuration file must be deployed each time a change is made in the fair scheduler configuration. The new file is read automatically by HiveServer2.

Placement Rule Requirement

After HiveServer2 determines into which queue the job can be placed, it submits the job with a queue specified in the submission request. This placement request is only honored if the placement rules in the `fair-scheduler.xml` file has the following rule defined as the first rule in the fair-scheduler policy:

```
<rule name="specified" />
```

Make sure that this is the first rule defined in the `fair-scheduler.xml` file. If another rule is defined as the first rule, placement behavior cannot be determined.

Job Submission Access Requirement

The `hive` user must be in the list of users allowed to submit jobs to the pool. The submitting user is used to determine pool placement, but the `hive` user is still the user that is expected to run the job in the pool.

Sentry does not manage permissions of Hive directories after certain commands

Sentry stops managing HDFS permissions of Hive directories that were the targets of certain DDL/DML commands.

Examples:

Alter tables, which are not rename or set location. For example "alter table set property"

Insert on a unpartitioned table.

Products affected: Sentry, HDFS, Hive

Affected Versions: All versions of CDH 5, except for those indicated in the 'Addressed in release/refresh/patch' section below.

Users affected: Users running Sentry with Hive or Impala using the HDFS ACL synchronization feature.

Impact:

The table/partition directory ACLs and permissions will be reverted to the underlying HDFS values after either of these workflows

Insert on unpartitioned tables

Alter table property on any table

Immediate action required: Upgrade to fixed version as soon as possible. If affected prior to an upgrade restart the HMS to act as a reset until the DDL/DML commands are rerun.

Addressed in release/refresh/patch: CDH 5.5.5 and higher, CDH 5.7.2 and higher, CDH 5.8.0 and higher

CREATE FUNCTION ... USING JAR does not work on Sentry-secured clusters

In a cluster without Sentry, a user is able to create a UDF using the `CREATE FUNCTION ... USING <hdfs location>` command in Hive, with a JAR located on HDFS. However, once Sentry is enabled, this command does not work even if the user is granted the `ALL` privilege to the URI on HDFS.

Affected Versions: CDH 5.7, 5.6, 5.5, 5.4

Fixed Versions: CDH 5.9.0 and above

Cloudera Issue: CDH-33816

With Sentry enabled, only Hive admin users have access to YARN job logs

As a prerequisite of enabling Sentry, Hive impersonation is turned off, which means all YARN jobs are submitted to the Hive job queue, and are run as the `hive` user. This is an issue because the YARN History Server now has to block users from accessing logs for their own jobs, since their own usernames are not associated with the jobs. As a result, end users cannot access any job logs unless they can get `sudo` access to the cluster as the `hdfs`, `hive` or other admin users.

In CDH 5.8 (and higher), Hive overrides the default configuration, `mapred.job.queue.name`, and places incoming jobs into the connected user's job queue, even though the submitting user remains `hive`. Hive obtains the relevant queue/username information for each job by using YARN's `fair-scheduler.xml` file.

Affected Versions: All CDH 5.x versions

Cloudera Issue: CDH-22890

Moving a partitioned table to a new location on the filesystem does not affect ACLs set on the previous location

With HDFS/Sentry sync enabled, if you move a partitioned table to a new location on the filesystem using the `ALTER TABLE .. SET LOCATION` command, ACLs set on the previous location remain unchanged. This occurs irrespective of whether the table is managed by Sentry.

Affected Versions: All CDH 5.x versions

Apache Issue: [SENTRY-1373](#)

Cloudera Issue: CDH-41784

Column-level privileges are not supported on Hive Metastore views

GRANT and REVOKE for column level privileges is not supported on Hive Metastore views.

Apache Issue: [SENTRY-754](#)

SELECT privilege on all columns does not equate to SELECT privilege on table

Users who have been explicitly granted the SELECT privilege on all columns of a table, will *not* have the permission to perform table-level operations. For example, operations such as SELECT COUNT (1) or SELECT COUNT (*) will not work even if you have the SELECT privilege on all columns.

There is one exception to this. The SELECT * FROM TABLE command will work even if you do not have explicit table-level access.

Apache Issue: [SENTRY-838](#)

The EXPLAIN SELECT operation works without table or column-level privileges

Users are able to run the EXPLAIN SELECT operation, exposing metadata for all columns, even for tables/columns to which they weren't explicitly granted access.

Apache Issue: [SENTRY-849](#)

With HDFS sync enabled, unexpected directory permissions are set when the NameNode plugin cannot communicate with the Sentry Server

With HDFS-Sentry sync enabled, if the NameNode plugin is unable to communicate with the Sentry Service, affected HDFS files will continue to use a cached copy of the synchronized ACLs for a configurable period of time, after which they will fall back to the Hive System User and the Hive System Group (for example, hive:hive). The timeout value can be modified by adding the `sentry.authorization-provider.cache-stale-threshold.ms` parameter to the `hdfs-site.xml` Safety Valve in Cloudera Manager. The default timeout value is 60 seconds, but you can increase this value from several minutes to a few hours, as needed to accommodate large clusters.

Affected Versions: CDH 5.14.7, 5.5.2

Fixed Versions: CDH 5.5.0 and above

Cloudera Issue: CDH-33623

Hive authorization (Grant/Revoke/Show) statements do not support fully qualified table names (`default.tab1`)

Workaround: Switch to the database before granting privileges on the table.

Affected Versions: CDH 5.12.x and below

Fixed Versions: CDH 5.13.0 and above

Cloudera Issue: CDH-19530

Object types Server and URI are not supported in SHOW GRANT ROLE `roleName` on OBJECT `objectName`

Workaround: Use SHOW GRANT ROLE `roleName` to list all privileges granted to the role.

Affected Versions: All CDH 5.x versions

Cloudera Issue: CDH-19430

Relative URI paths not supported by Sentry

Sentry supports only absolute (not relative) URI paths in permission grants. Although some early releases (for example, CDH 5.7.0) may not have raised explicit errors when relative paths were set, upgrading a system that uses relative paths causes the system to lose Sentry permissions.

Affected Versions: All versions. Relative paths are not supported in Sentry for permission grants.

Resolution: Revoke privileges that have been set using relative paths, and grant permissions using absolute paths before upgrading.

Absolute (Use this form)	Relative (Do not use this form)
hdfs://absolute/path/	hdfs://relative/path
s3a://bucketname/	s3a://bucketname

Apache Spark Known Issues

CVE-2019-10099: Apache Spark local files left unencrypted

Certain operations in Spark leave local files unencrypted on disk, even when local file encryption is enabled with “spark.io.encryption.enabled”.

This includes cached blocks that are fetched to disk (controlled by `spark.maxRemoteBlockSizeFetchToMem`) in the following cases:

- In SparkR when `parallelize` is used
- In Pyspark when `broadcast` and `parallelize` are used
- In Pyspark when `python udfs` is used

Products affected:

- CDH
- CDS Powered by Apache Spark

Affected versions:

- CDH 5.15.1 and earlier
- CDH 6.0.0
- CDS 2.1.0 release 1 and release 2
- CDS 2.2.0 release 1 and release 2
- CDS 2.3.0 release 3

Users affected: All users who run Spark on CDH and CDS in a multi-user environment.

Date/time of detection: July 2018

Severity (Low/Medium/High): 6.3 Medium ([CVSS AV:L/AC:H/PR:N/UI:R/S:U/C:H/I:H/A:N](#))

Impact: Unencrypted data accessible.

CVE: CVE-2019-10099

Immediate action required: Upgrade to a version of CDH containing the fix.

Workaround: Do not use of pyspark and the fetch-to-disk options.

Fixed versions:

- CDH 5.15.2
- CDH 5.16.0
- CDH 6.0.1
- CDS 2.1.0 release 3
- CDS 2.2.0 release 3
- CDS 2.3.0 release 4

For the latest update on this issue see the corresponding Knowledge article: [TSB 20210-336: Apache Spark local files left unencrypted](#)

Partitioned tables in the Avro format with upper case column names return NULL values from Spark SQL

When you use Spark SQL to query external partitioned Hive tables created in the Avro format and which contain upper case column names, Spark SQL returns NULL values for the upper case column names.

Workaround: In Spark 1.6, create aliases that do not contain upper case characters for each column name.

Affected Versions: CDH 5.8.2 - CDH 5.14.3

Fixed Versions: CDH 5.14.4 and later

Apache Bug: [SPARK-15848](#)

Cloudera Bug: CDH-72396, CDH-45978

Missing results in Hive, Spark, Pig, Custom MapReduce jobs, and other Java applications when filtering Parquet data written by Impala

Apache Hive and Apache Spark rely on Apache Parquet's `parquet-mr` Java library to perform filtering of Parquet data stored in row groups. Those row groups contain statistics that make the filtering efficient without having to examine every value within the row group.

Recent versions of the `parquet-mr` library contain a bug described in [PARQUET-1217](#). This bug causes filtering to behave incorrectly if only some of the statistics for a row group are written. Starting in CDH 5.13, Apache Impala populates statistics in this way for Parquet files. As a result, Hive and Spark may incorrectly filter Parquet data that is written by Impala.

In CDH 5.13, Impala started writing Parquet's `null_count` metadata field without writing the `min` and `max` fields. This is valid, but it triggers the [PARQUET-1217](#) bug in the predicate push-down code of the Parquet Java library (`parquet-mr`). If the `null_count` field is set to a non-zero value, `parquet-mr` assumes that `min` and `max` are also set and reads them without checking whether they are actually there. If those fields are not set, `parquet-mr` reads their default value instead.

For integer SQL types, the default value is 0, so `parquet-mr` incorrectly assumes that the `min` and `max` values are both 0. This causes the problem when filtering data. Unless the value 0 itself matches the search condition, all row groups are discarded due to the incorrect `min/max` values, which leads to missing results.

Affected Products: The Parquet Java library (`parquet-mr`) and by extension, all Java applications reading Parquet files, including, but not limited to:

- Hive
- Spark
- Pig
- Custom MapReduce jobs

Affected Versions:

- CDH 5.13.0, 5.13.1, 5.13.2, and 5.14.0
- CDS 2.2 Release 2 Powered by Apache Spark and earlier releases on CDH 5.13.0 and later

Who Is Affected: Anyone writing Parquet files with Impala and reading them back with Hive, Spark, or other Java-based components that use the `parquet-mr` libraries for reading Parquet files.

Severity (Low/Medium/High): High

Impact: Parquet files containing null values for integer fields written by Impala produce missing results in Hive, Spark, and other Java applications when filtering by the integer field.

Immediate Action Required:

- **Upgrade**

You should upgrade to one of the fixed maintenance releases mentioned below.

- **Workaround**

This issue can be avoided at the price of performance by disabling predicate push-down optimizations:

- In Hive, use the following `SET` command:

```
SET hive.optimize.ppd = false;
```
- In Spark, disable the following configuration setting:

```
--conf spark.sql.parquet.filterPushdown=false
```

Addressed in the Following Releases:

- CDH 5.13.3 and higher
- CDH 5.14.2 and higher
- CDH 5.15.0 and higher
- CDS 2.3 Release 2 and higher

For the latest update on this issue, see the corresponding Knowledge Base article:

[TSB:2018-300: Missing results in Hive, Spark, Pig, and other Java applications when filtering Parquet data written by Impala](#)

Apache Spark experimental features are not supported unless specifically identified as supported

If an Apache Spark feature or API is identified as experimental, in general Cloudera does not provide support for it.

Spark 2.x not supported on unmanaged clusters

Launching Spark 2.x jobs from machines not managed by Cloudera Manager is not supported.

ADLS not Supported for All Spark Components

Microsoft Azure Data Lake Store (ADLS) is a cloud-based filesystem that you can access through Spark applications. Spark with Kudu is not currently supported for ADLS data. (Hive-on-Spark is available for ADLS in CDH 5.12 and higher.)

IPython / Jupyter notebooks not supported

- The IPython notebook system (renamed to Jupyter as of IPython 4.0) is not supported.

Certain Spark Streaming features not supported

- The `mapWithState` method is unsupported because it is a nascent unstable API.
- **Affected Versions:** All CDH 5 versions
- **Bug:** [SPARK-2629](#).

Certain Spark SQL features not supported

The following Spark SQL features are not supported:

- Thrift JDBC/ODBC server
- Spark SQL CLI

Spark Dataset API not supported

Cloudera distribution of Spark 1.6 does not support the Spark Dataset API. However, Spark 2.0 and higher supports the Spark Dataset API.

JDBC Datasource API not supported

Using the JDBC Datasource API to access Hive or Impala is not supported

GraphX not supported

Cloudera does not support GraphX.

Affected Versions: All CDH 5 versions.

SparkR not supported

Cloudera does not support SparkR.

Affected Versions: All CDH 5 versions.

Scala 2.11 not supported

Spark in CDH 5.x does not support Scala 2.11 because it is binary incompatible, and also not yet full-featured. However, Spark 2.0 and higher supports Scala 2.11.

Affected Versions: All CDH 5 versions.

Spark Streaming cannot consume from secure Kafka till it starts using Kafka 0.9 Consumer API

Bug: [SPARK-12177](#).

Workaround: Use Spark 2.

Tables saved with the Spark SQL `DataFrame.saveAsTable` method are not compatible with Hive

Writing a `DataFrame` directly to a Hive table creates a table that is not compatible with Hive; the metadata stored in the metastore can only be correctly interpreted by Spark. For example:

```
val hsc = new HiveContext(sc)
import hsc.implicits._
val df = sc.parallelize(data).toDF()
df.write.format("parquet").saveAsTable(tableName)
```

creates a table with this metadata:

```
inputFormat:org.apache.hadoop.mapred.SequenceFileInputFormat
outputFormat:org.apache.hadoop.hive.ql.io.HiveSequenceFileOutputFormat
```

This is also occurs when using explicit schema, such as:

```
val schema = StructType(Seq(...))
val data = sc.parallelize(Seq(Row(...), ...))
val df = hsc.createDataFrame(data, schema)
df.write.format("parquet").saveAsTable(tableName)
```

Affected Versions: CDH 5.5 and higher

Workaround: Explicitly create a Hive table to store the data. For example:

```
df.registerTempTable(tempName)
hsc.sql(s"""
CREATE TABLE $tableName (
// field definitions )
STORED AS $format """)
hsc.sql(s"INSERT INTO TABLE $tableName SELECT * FROM $tempName")
```

Cloudera Bug: CDH-33639

Cannot create Parquet tables containing date fields in Spark SQL

If you create a Parquet table containing date field in Spark SQL, you see the following exception:

```
Exception in thread "main" org.apache.spark.sql.execution.QueryExecutionException:
FAILED: Execution Error, return code 1 from
org.apache.hadoop.hive.ql.exec.DDLTask.java.lang.UnsupportedOperationException: Parquet
does not support date.
at
org.apache.spark.sql.hive.client.ClientWrapper$$anonfun$runHive$1.apply(ClientWrapper.scala:433)
```

This is due to a limitation ([HIVE-6384](#)) in the version of Hive (1.1) included in CDH 5.5.0.

Cloudera Bug: CDH-33640

Spark SQL does not support the union type

Tables containing union fields cannot be read or created using Spark SQL.

Affected Versions: CDH 5.5.0 and higher

Cloudera Bug: CDH-33641

Spark SQL does not respect size limit for the varchar type

Spark SQL treats `varchar` as string (that is, there no size limit). The observed behavior is that Spark reads and writes these columns as regular strings; if inserted values exceed the size limit, no error will occur. The data will be truncated when read from Hive, but not when read from Spark.

Affected Versions: CDH 5.5.0 and higher

Bug: [SPARK-5918](#)

Cloudera Bug: CDH-33642

Spark SQL does not support the char type

Spark SQL does not support the `char` type (fixed-length strings). Like unions, tables with such fields cannot be created from or read by Spark.

Affected Versions: CDH 5.5.0 to CDH 5.6.1

Fixed in Versions: CDH 5.7.0 and higher

Cloudera Bug: CDH-33643

Spark SQL does not support transactional tables

Spark SQL does not support Hive transactions ("ACID").

Affected Versions: CDH 5.5.0 and higher

Cloudera Bug: CDH-33644

Spark SQL does not prevent you from writing key types not supported by Avro tables

Spark allows you to declare DataFrames with any key type. Avro supports only string keys and trying to write any other key type to an Avro table will fail.

Affected Versions: CDH 5.5.0 and higher

Cloudera Bug: CDH-33648

Spark SQL does not support timestamp in Avro tables

Affected Versions: CDH 5.5.0 and higher

Cloudera Bug: CDH-33649

Spark SQL does not support all 'ANALYZE TABLE COMPUTE STATISTICS' syntax

`ANALYZE TABLE <table name> COMPUTE STATISTICS NOSCAN` works. `ANALYZE TABLE <table name> COMPUTE STATISTICS (without noscan)` and `ANALYZE TABLE <table name> COMPUTE STATISTICS FOR COLUMNS` both return errors.

Affected Versions: CDH 5.5.0 and higher

Cloudera Bug: CDH-33650

Spark SQL statements that can result in table partition metadata changes may fail

Because Spark does not have access to Sentry data, it may not know that a user has permissions to execute an operation and instead fail it. SQL statements that can result in table partition metadata changes, for example, "ALTER TABLE" or "INSERT", may fail.

Affected Versions: CDH 5.5.0 to CDH 5.6.1

Fixed in Versions: 5.7.0 and higher

Cloudera Bug: CDH-33446

Spark SQL does not respect Sentry ACLs when communicating with Hive metastore

Even if user is configured via Sentry to not have read permission to a Hive table, a Spark SQL job running as that user can still read the table's metadata directly from the Hive metastore.

Cloudera Bug: CDH-33658

Dynamic allocation and Spark Streaming

If you are using Spark Streaming, Cloudera recommends that you disable dynamic allocation by setting `spark.dynamicAllocation.enabled` to `false` when running streaming applications.

Spark uses Akka version 2.2.3

The CDH 5.5 version of Spark 1.5 differs from the Apache Spark 1.5 release in using Akka version 2.2.3, the version used by Spark 1.1 and CDH 5.2. Apache Spark 1.5 uses Akka version 2.3.11.

Spark standalone mode does not work on secure clusters

Cloudera Bug: CDH-16998, "fixed" by workaround.

Workaround: On secure clusters, run Spark applications on YARN.

Limitation with Region Pruning for HBase Tables

When SparkSQL accesses an HBase table through the HiveContext, region pruning is not performed. This limitation can result in slower performance for some SparkSQL queries against tables that use the HBase SerDes than when the same table is accessed through Impala or Hive.

Input Format Change on `spark.yarn.am.waitTime`

The `spark.yarn.am.waitTime` property does not accept valid input formats, such as "250s".

Affected Versions: CDH 5.10.0 to CDH 5.10.2

Fixed in Versions: CDH 5.10.3, CDH 5.11.3, CDH 5.12.2, CDH 5.13.0 and higher

Workaround: Specify the value in milliseconds without "s" at the end.

History link in ResourceManager web UI broken for killed Spark applications

When a Spark application is killed, the history link in the ResourceManager web UI does not work.

Workaround: To view the history for a killed Spark application, see the Spark HistoryServer web UI instead.

Affected Versions: All CDH versions

Apache Issue: None

Cloudera Issue: CDH-49165

Spark application using dynamic allocation fails with `java.lang.AssertionError`

When dynamic allocation is enabled, Spark applications can fail with a `java.lang.AssertionError: assertion failed` error when the number of running executors exceeds the target number of executors. For example, this can

occur if an application has many small tasks, and Spark requests a large amount of executors before reducing the target, triggering a race condition resulting in an error message similar to the following

```
ERROR yarn.YarnAllocator: Failed to launch executor 16 on container
container_e11_1513498336069_33453_01_000022
java.lang.AssertionError: assertion failed
```

Workaround: Disable dynamic allocation by specifying `spark.dynamicAllocation.enabled=false` for the problematic job.

Affected Versions:

- CDH versions lower than 5.12
- CDH 5.12.0, 5.12.1
- CDH 5.13.0

Fixed Versions:

- CDH 5.12.2
- CDH 5.13.1
- CDH 5.14 and higher

Apache Issue: [SPARK-17511](#)

Cloudera Issue: CDH-56799

Apache Sqoop Known Issues

MySQL JDBC driver shipped with CentOS 6 systems does not work with Sqoop

CentOS 6 systems currently ship with version 5.1.17 of the MySQL JDBC driver. This version does not work correctly with Sqoop.

Bug: None

Affected Versions: MySQL JDBC 5.1.17, 5.1.4, 5.3.0

Cloudera Bug: CDH-23180

Workaround: Install version 5.1.31 of the JDBC driver as detailed in (Sqoop 1) or (Sqoop 2).

MS SQL Server "integratedSecurity" option unavailable in Sqoop

The `integratedSecurity` option is not available in the Sqoop CLI.

Bug: None

Workaround: None

Sqoop 1

Sqoop does not support Hive metastore HA for Parquet format

If the Hive metastore (HMS) is configured for high availability (HA), the value of the `hive.metastore.uris` property in `/etc/hive/conf.cloudera.HIVE-1/hive-site.xml` contains a list of URIs instead of one URI. For example:

```
<value>thrift://hive-1.example.com:9083,thrift://hive-2.example.com:9083,thrift://hive-3.example.com:9083</value>
```

Sqoop uses Kite for Parquet file handling and Kite does not support HMS HA. Using Kite, Sqoop takes the first URI from the `hive.metastore.uris` list and uses that to connect to the HMS. If that address is not available, the Sqoop command fails.

This issue occurs when you use the following Sqoop arguments with the `sqoop-import` command: `--hive-import` and `--as-parquetfile`.

Bug: CDH-57895

Affected Versions: CDH 5.12.x and below.

Fixed in Versions: CDH 5.13 and higher.

Cloudera Bug: CDH-57895

Workaround: None.

Sqoop1 Metastore job tool ignores overrides

Sqoop job command option overrides are being ignored due to SQOOP-2896. The `exec` action should let you override arguments in a saved job by passing `options` after `--`, as in:

```
sqoop job --exec job -- options_to_override>
```

However, the values in `options_to_override` are not being passed to the saved job. Sqoop jobs that use overrides can break or execute incorrectly, as shown in the example below.

```
sqoop job --create my_eval_query -- eval --connect jdbc:mysql://fqdn.example.com:3306/db
--username user --password password --query "select * from t1"
```

This command executes the job defined in `my_eval_query` and returns the rows in the table:

```
sqoop job --exec my_eval_query
c1  c2          c3
1   2017-02-17 row data 1
2   2017-02-17 row data 2
3   2017-02-17 row data 3
```

In this next run below, the override query is ignored. The job uses the original parameters stored for the job and returns the same rows as before rather than the single row where column `c1 = 2`:

```
sqoop job --exec my_eval_query -- --query "select * from t1 where c1 = 2"
c1  c2          c3
1   2017-02-17 row data 1
2   2017-02-17 row data 2 <--- Only this row should be returned
3   2017-02-17 row data 3
```

Affected releases: CDH 5.8.0, 5.8.1, 5.8.3, 5.8.4; CDH 5.9.0, 5.9.1; CDH 5.10.0

Bug: SQOOP-2896

Cloudera Bug: CDH-51101

Fixed in Versions: CDH 5.8.5, CDH 5.9.2, CDH 5.9.3, CDH 5.10.1, CDH 5.10.2, CDH 5.11.0 (and higher) releases.

Sqoop1 (doc import + --as-parquetfile) limitation with KMS/KTS Encryption at Rest

Due to a limitation with Kite SDK, it is not possible to use (sqoop import --as-parquetfile) with KMS/KTS Encryption zones. See the following example.

```
sqoop import --connect jdbc:db2://djaxludb1001:61035/DBAT003 --username=dh810202 --P
--target-dir /data/hive_scratch/ASDISBURSEMENT --delete-target-dir -m1 --query "select
disbursementnumber,disbursementdate,xmldata FROM DB2dba.ASDISBURSEMENT where
DISBURSEMENTNUMBER = 2011113210000115311 AND \${CONDITIONS}" -hive-import --hive-database
adminserver -hive-table asdisbursement_dave --map-column-java XMLDATA=String
--as-parquetfile
```

```
16/12/05 12:23:46 INFO mapreduce.Job: map 100% reduce 0%
16/12/05 12:23:46 INFO mapreduce.Job: Job job_1480530522947_0096 failed with state FAILED
due to: Job commit failed: org.kitesdk.data.DatasetIOException: Could not move contents
of
hdfs://AJAX01-ns/tmp/adminserver/.temp/job_1480530522947_0096/mr/job_1480530522947_0096
to hdfs://AJAX01-ns/data/RetiredApps/INS/AdminServer/asdisbursement_dave
<SNIP>
Caused by: org.apache.hadoop.ipc.RemoteException(java.io.IOException):
```

```
/tmp/adminserver/.temp/job_1480530522947_0096/mr/job_1480530522947_0096/5dbcac42-5d69-4e46-88c2-171bedac4858.parquet
can't be moved into an encryption zone.
```

Cloudera Bug: CDH-40826 and SQOOP-2943

Workaround: Use an alternate data file type, for example text or avro.

[Doc import as Parquet files may result in out-of-memory errors](#)

Out-of-memory (OOM) errors can be caused in the following two cases:

- With many very large rows (multiple megabytes per row) before initial-page-run check (ColumnWriter)
- When rows vary significantly by size so that the next-page-size check is based on small rows and is set very high followed by many large rows

Bug: PARQUET-99

Cloudera Bug: CDH-50896

Workaround: None, other than restructuring the data.

[Hive, Pig, and Sqoop 1 fail in MRv1 tarball due to incorrect HADOOP_MAPRED_HOME setting](#)

This issue affects tarball installations only.

Bug: None

Cloudera Bug: CDH-6640. Fixed for packages only: CDH-6648.

Workaround: For MRv1 only, manually edit as follows:

- Change the path for the Hadoop MapReduce home setting in `/etc/default/hadoop` from:

```
export HADOOP_MAPRED_HOME=/usr/lib/hadoop-mapreduce
```

to:

```
export HADOOP_MAPRED_HOME=/usr/lib/hadoop-0.20-mapreduce
```

- Remove `/usr/lib/hadoop-mapreduce` from the `HADOOP_CLASSPATH`.

[Sqoop 2](#)

[Sqoop 2 client cannot be used with a different version of the Sqoop 2 server](#)

The Sqoop 2 client and server must be running the same CDH version.

Cloudera Bug: CDH-11901

Workaround: Make sure all Sqoop 2 components are running the same version of CDH.

[Sqoop 2 upgrade may fail if any job's source and destination links point to the same connector](#)

For example, the links for the job shown in the following output both point to `generic-jdbc-connector`:

```
sqoop:000> show job --all
1 job(s) to show:
Job with id 1 and name job1 (Enabled: true, Created by null at 5/13/15 3:05 PM, Updated
by null at 5/13/15 6:04 PM)
  Throttling resources
    Extractors:
    Loaders:
From link: 1
  From database configuration
    Schema name: schemal
    Table name: tabl
    Table SQL statement:
    Table column names: coll
    Partition column name:
```

```

Null value allowed for the partition column: false
Boundary query:
Incremental read
Check column:
Last value:
To link: 2
To database configuration
Schema name: schema2
Table name: tab2
Table SQL statement:
Table column names: col2
Stage table name:
Should clear stage table:

sqoop:000> show link --all
2 link(s) to show:
link with id 1 and name try1 (Enabled: true, Created by null at 5/13/15 2:59 PM, Updated
by null at 5/13/15 5:47 PM)
Using Connector generic-jdbc-connector with id 2
Link configuration
JDBC Driver Class: com.mysql.jdbc.Driver
JDBC Connection String: jdbc:mysql://mysql.server/database
Username: nvaidya
Password:
JDBC Connection Properties:
link with id 2 and name try2 (Enabled: true, Created by null at 5/13/15 3:01 PM, Updated
by null at 5/13/15 5:47 PM)
Using Connector generic-jdbc-connector with id 2
Link configuration
JDBC Driver Class: com.mysql.jdbc.Driver
JDBC Connection String: jdbc:mysql://mysql.server/database
Username: nvaidya
Password:
JDBC Connection Properties:

```

Cloudera Bug: CDH-27779

Workaround: Before upgrading, make sure no jobs have source and destination links that point to the same connector.

Apache Phoenix Known Issues

Potential deadlock on region opening with secondary index recovery

Distributed deadlock happens in clusters with a moderate number of regions for the data tables, secondary index tables.

Products affected: Phoenix

Releases affected: Phoenix 4.14.1 parcel

User affected: Users of Phoenix with secondary indexes

Severity (Low/Medium/High): Medium

Impact: Deadlock on region opening

Immediate action required/workaround:

Make the following configuration changes in HBase:

- 1) Set `hbase.master.startup.retainassign` to `false`
- 2) Increase the value of `hbase.regionserver.executor.openregion.threads` and restart cluster should bring up cluster.

If you continue to encounter issues, tune the following assignment manager parameters to match the count of regions for faster assignments:

- `hbase.assignment.threads.max`
- `hbase.master.namespace.init.timeout`

- `hbase.master.wait.on.regionserver.mintostart`
- `hbase.bulk.assignment.threshold.regions`
- `hbase.bulk.assignment.threshold.servers`

Apache ZooKeeper Known Issues

ZooKeeper JMX did not support TLS when managed by Cloudera Manager

The ZooKeeper service optionally exposes a JMX port used for reporting and metrics. By default, Cloudera Manager enables this port, but prior to Cloudera Manager 6.1.0, it did not support mutual TLS authentication on this connection. While JMX has a password-based authentication mechanism that Cloudera Manager enables by default, weaknesses have been found in the authentication mechanism, and Oracle now advises JMX connections to enable mutual TLS authentication in addition to password-based authentication. A successful attack may leak data, cause denial of service, or even allow arbitrary code execution on the Java process that exposes a JMX port. Beginning in Cloudera Manager 6.1.0, it is possible to configure mutual TLS authentication on ZooKeeper's JMX port.

Products affected: ZooKeeper

Releases affected: Cloudera Manager 6.1.0 and lower, Cloudera Manager 5.16 and lower

Users affected: All

Date/time of detection: June 7, 2018

Severity (Low/Medium/High): 9.8 High (CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H)

Impact: Remote code execution

CVE: CVE-2018-11744

Immediate action required: Upgrade to Cloudera Manager 6.1.0 and enable TLS for the ZooKeeper JMX port by turning on the configuration settings "Enable TLS/SSL for ZooKeeper JMX" and "Enable TLS client authentication for JMX port" on the ZooKeeper service and configuring the appropriate TLS settings. Alternatively, disable the ZooKeeper JMX port via the configuration setting "Enable JMX Agent" on the ZooKeeper service.



Note: Disabling the ZooKeeper JMX port prevents Cloudera Manager from performing health checks on the ZooKeeper service.

Addressed in release/refresh/patch: Cloudera Manager 6.1.0

Adding New ZooKeeper Servers Can Lead to Data Loss

When the number of new ZooKeeper servers exceeds the number that already exist in the ZooKeeper service (for example, if you increase the number of servers from 1 to 3), and a Start command is immediately issued to the ZooKeeper service, the new servers can form a quorum, which causes data loss in existing servers.

Users of the following versions of Cloudera Manager are affected:

5.0.0–5.0.5, 5.1.0–5.1.4, 5.2.0–5.2.4, and 5.3.0–5.3.2

Workaround: If you use a version of Cloudera Manager listed above, upgrade to the next available maintenance release with the bug fix (within the minor version), or to Cloudera Manager 5.4.

The ZooKeeper server cannot be migrated from version 3.4 to 3.3, then back to 3.4, without user intervention.

Upgrading from 3.3 to 3.4 is supported, as is downgrading from 3.4 to 3.3. However, moving from 3.4 to 3.3 and back to 3.4 will fail. 3.4 is checking the `datadir` for `acceptedEpoch` and `currentEpoch` files and comparing these against the snapshot and log files contained in the same directory. These epoch files are new in 3.4.

As a result: 1) Upgrading from 3.3 to 3.4 is fine - the `*Epoch` files do not exist, and the server creates them. 2) Downgrading from 3.4 to 3.3 is also fine as version 3.3 ignores the `*Epoch` files. 3) Going from 3.4 to 3.3 then back to 3.4 fails because 3.4 sees invalid `*Epoch` files in the `datadir`; 3.3 will have ignored them, applying changes to the snapshot and log files without updating the `*Epoch` files.

Bug: [ZOOKEEPER-1149](#)

Cloudera Bug: CDH-5272

Anticipated Resolution: See workaround

Workaround: Delete the *Epoch files if this situation occurs — the version 3.4 server will recreate them as in case 1) above.

Issues Fixed in CDH 5



Note: For links to the detailed change lists that describe the bug fixes and improvements to all of the CDH 5 projects, including bug-fix reports for the corresponding upstream Apache projects, see the packaging section of [CDH Version and Packaging Information](#).

Issues Fixed in CDH 5.16.x

The following topics describe issues fixed in CDH 5.16.x, from newest to oldest release. You can also review [What's New in Cloudera Documentation](#) or [Known issues in CDH 5](#) for more information. Note that there was no CDH release numbered 5.16.0.

Issues Fixed in CDH 5.16.2

CDH 5.16.2 fixes the following issues:

XSS Cloudera Manager

Malicious Impala queries can result in Cross Site Scripting (XSS) when viewed in Cloudera Manager.

Products affected: Apache Impala

Releases affected:

- Cloudera Manager 5.13.x, 5.14.x, 5.15.1, 5.15.2, 5.16.1
- Cloudera Manager 6.0.0, 6.0.1, 6.1.0

Users affected: All Cloudera Manager Users

Date/time of detection: November 2018

Severity (Low/Medium/High): High

Impact: When a malicious user generates a piece of JavaScript in the impala-shell and then goes to the Queries tab of the Impala service in Cloudera Manager, that piece of JavaScript code gets evaluated, resulting in an XSS.

CVE: CVE-2019-14449

Immediate action required: There is no workaround, upgrade to the latest available maintenance release.

Addressed in release/refresh/patch:

- Cloudera Manager 5.16.2
- Cloudera Manager 6.0.2, 6.1.1, 6.2.0, 6.3.0

Timestamp type-casted to varchar in a binary predicate can produce incorrect result

In an Impala query the timestamp can be type-casted to a varchar of smaller length to convert a timestamp value to a date string. However, if such Impala query is used in a binary comparison against a string literal, it can produce incorrect results, because of a bug in the expression rewriting code. The following is an example of this:

```
> select * from (select cast('2018-12-11 09:59:37' as timestamp) as ts) tbl where cast(ts as varchar(10)) = '2018-12-11';
```

The output will have 0 rows.

Products affected: Apache Impala

Releases affected:

- CDH 5.15.0, 5.15.1, 5.15.2, 5.16.0, 5.16.1
- CDH 6.0.0, 6.0.1, 6.1.0, 6.1.1

For the latest update on this issue see the corresponding Knowledge article: [TSB 2019-358: Timestamp type-casted to varchar in a binary predicate can produce incorrect result](#)

Inconsistent rows returned from queries in Kudu

Due to KUDU-2463, upon restarting Kudu, inconsistent rows may be returned from tables that have not recently been written to, resulting in any of the following:

- multiple rows for the same key being returned
- deleted data being returned
- inconsistent results consistently being returned for the same query

If this happens, you have two options to resolve the conflicts: write to the affected Kudu partitions by:

- re-deleting the known and deleted data
- upserting the most up-to-date version of affected rows.

Products affected: Apache Kudu

Releases affected:

- CDH 5.12.2, 5.13.3, 5.14.4, 5.15.1, 5.16.1
- CDH 6.0.1, 6.1.0, 6.1.1

For the latest update on this issue see the corresponding Knowledge article: [TSB 2019-353: Inconsistent rows returned from queries in Kudu](#)

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.16.2:

Apache Flume

- [FLUME-2973](#) - Deadlock in hdfs sink.
- [FLUME-3223](#) - Flume HDFS Sink should retry close prior recover lease.

Apache Hadoop

- [HADOOP-15442](#) - ITestS3AMetrics.testMetricsRegister should not know the name of the metrics source.
- [HDFS-11751](#) - DFSZKFailoverController daemon exits with the wrong status code.
- [HDFS-12683](#) - DFSZKFailOverController re-order logic for logging exception.
- [HDFS-14111](#) - hdfsOpenFile on HDFS causes unnecessary IO from file offset 0
- [MAPREDUCE-6382](#) - HTML links in the Diagnostics in JHS job overview must not be escaped.
- [MAPREDUCE-7125](#) - JobResourceUploader creates LocalFileSystem when it's not necessary.
- [MAPREDUCE-7131](#) - Job History Server has race condition where it moves files from intermediate to finished but thinks file is in intermediate.
- [YARN-4227](#) - Ignore expired containers from the removed nodes in FairScheduler.
- [YARN-4677](#) - RMNodeResourceUpdateEvent update from scheduler can lead to race condition.

Apache HBase

- [HBASE-16810](#) - HBase Balancer throws ArrayIndexOutOfBoundsException when regionservers are in /hbase/draining znode and unloaded
- [HBASE-17510](#) - DefaultMemStore gets the wrong heap size after rollback
- [HBASE-19730](#) - Backport HBASE-14497 Reverse Scan threw StackOverflow caused by readPt checking
- [HBASE-20604](#) - ProtobufLogReader#readNext can incorrectly loop to the same position in the stream until the the WAL is rolled

- [HBASE-21275](#) - Disable TRACE HTTP method for thrift http server
- [HBASE-21546](#) - ConnectException in TestThriftHttpServer

Apache Hive

- [HIVE-12476](#) - Metastore NPE on Oracle with Direct SQL
- [HIVE-13278](#) - Avoid FileNotFoundException when map/reduce.xml is not available
- [HIVE-13394](#) - Analyze table fails in Tez on empty partitions
- [HIVE-13592](#) - metastore calls map is not thread safe
- [HIVE-14557](#) - Nullpointer When both SkewJoin and Mapjoin Enabled
- [HIVE-14560](#) - Support exchange partition between S3 and HDFS tables
- [HIVE-14690](#) - Query fail when hive.exec.parallel=true, with conflicting session dir
- [HIVE-16839](#) - Unbalanced calls to openTransaction/commitTransaction when altering the same partition concurrently
- [HIVE-18778](#) - Needs to capture input/output entities in explain
- [HIVE-20331](#) - Query with union all, lateral view and Join fails with "cannot find parent in the child operator"
- [HIVE-20678](#) - HiveHBaseTableOutputFormat should implement HiveOutputFormat to ensure compatibility
- [HIVE-20695](#) - HoS Query fails with hive.exec.parallel=true
- [HIVE-21028](#) - get_table_meta should use a fetch plan to avoid race conditions ending up in NucleusObjectNotFoundException
- [HIVE-21044](#) - Improvements to HMS metrics
- [HIVE-21045](#) - Add connection pool info and rolling performance info to the metrics system

Hue

- [HUE-8388](#) - [oozie] Make Hue create a new workspace when importing an Oozie workflow instead of using the "deployment_dir" field
- [HUE-8450](#) - [editor] Embedded mode improvements for previous Hue version
- [HUE-8458](#) - [frontend] Improve application loading performance
- [HUE-8468](#) - [frontend] Dynamically adding styles in embedded mode fails in Internet Explorer (throws a Java script exception)
- [HUE-8584](#) - [useradmin] Errors returned for Add Sync Ldap Group
- [HUE-8585](#) - [useradmin] Errors returned for Add Sync Ldap Users
- [HUE-8631](#) - HBase is not accessible by way of the Hue server; instead returns "API Error."
- [HUE-8660](#) - [core] File browser cannot view files containing a hash (#) in the name
- [HUE-8691](#) - [useradmin] Attempting to add/sync group will not add users if the objectClass posixGroup exists in the LDAP entry
- [HUE-8692](#) - [useradmin] Group sync fails if all group members are not found with error "No such object"
- [HUE-8693](#) - [useradmin] Security application only displays 100 users in the impersonation list
- [HUE-8705](#) - [oozie] Hidden popup window is blocking the Query drop-down menu and the search box
- [HUE-8709](#) - [useradmin] Black transparent screen remains after confirmation modal is hidden
- [HUE-8746](#) - [pig] Add hcat support to the Pig Editor in Hue

Hue allows unsigned SAML assertions

If Hue receives an unsigned assertion, it continues to process it as valid. This means it is possible for an end-user to forge or remove the signature and manipulate a SAML assertion to gain access without a successful authentication.

Products affected: Hue, CDH

Releases affected:

- CDH 5.15.x and earlier
- CDH 5.16.0, 5.16.1
- CDH 6.0.x
- CDH 6.1.x

User affected: All users who are using SAML with Hue.

CVE: CVE-2019-14775

Date/time of detection: January 2019

Detected by: Joel Snape

Severity (Low/Medium/High): High

Impact:

This is a significant security risk as it allows anyone to fake their access validity and therefore access Hue, even if they should not have access. In more detail: if Hue receives an unsigned assertion, it continues to process it as valid. This means it is possible for an end-user to forge or remove the signature and manipulate a SAML assertion to gain access without a successful authentication.

CVE: CVE-2019-14775

Immediate action required:

- **Upgrade (recommended):** Upgrade to a version of CDH containing the fix.
- **Workaround:** None

Addressed in release/refresh/patch:

- CDH 5.16.2
- CDH 6.2.0

Apache Impala

- [IMPALA-6323](#) - Impala now supports a constant in the window specifications.
- [IMPALA-7960](#) - Impala now returns a correct result when comparing `TIMESTAMP` to a string literal in a binary predicate where the `TIMESTAMP` is casted to `VARCHAR` of smaller length.
- [IMPALA-7961](#) - Fixed an issue where queries running with the `SYNC_DDL` query option can fail when the Catalog Server is under a heavy load with concurrent catalog operations of long-running DDLs.
- [IMPALA-8058](#) - Fixed cardinality estimates for HBase queries, which could sometimes yield hugely high numbers.
- [IMPALA-8109](#) - Impala can now read the gzip files bigger than 2 GB.
- [IMPALA-8212](#) - Fixed a race condition in the Kerberos authentication code.

Kite

- [KITE-1185](#) - Make root temp directory path configurable in `HiveAbstractDatasetRepository`

Apache Kudu

- [KUDU-1678](#) - Fixed a crash caused by a race condition between altering tablet schemas and deleting tablet replicas.
- [KUDU-2195](#) - Now you can use the `--cmeta_force_fsync` flag to `fsync` Kudu's consensus metadata more aggressively. Setting this to `true` may decrease Kudu's performance, but will improve its durability in the face of power failures and forced shutdowns. The issue was much more likely to happen when Kudu was running on XFS.
- [KUDU-2463](#) - Fixed an issue in which incorrect results would be occasionally returned in scans following a server restart.
- [CDH-76920](#) - Fixed the issue where the Kudu CLI crashes when running the `'kudu cluster rebalance'` sub-command on some platforms.

Apache Oozie

- [OOZIE-3382](#) - Implement and backport `Optimize SshActionExecutor's drainBuffers` method

Apache Pig

- [PIG-5373](#) - `InterRecordReader` might skip records if certain sync markers are used
- [PIG-5374](#) - Use `CircularFifoBuffer` in `InterRecordReader`

Apache Sentry

- [SENTRY-2205](#) - Improve Sentry NN Logging.
- [SENTRY-2301](#) - Log where sentry stands in the snapshot fetching process, periodically

- [SENTRY-2372](#) - SentryStore should not implement grantOptionCheck
- [SENTRY-2419](#) - Log where sentry stands in the process of persisting the snapshot
- [SENTRY-2427](#) - PortUse Hadoop KerberosName class to derive shortName
- [SENTRY-2428](#) - Skip null partitions or partitions with null sds entries
- [SENTRY-2437](#) - PortWhen granting privileges a single transaction per grant causes long delays
- [SENTRY-2490](#) - PortWhen building a full perm update for each object we only build 1 privilege per role
- [SENTRY-2498](#) - Exception while deleting paths that does't exist
- [SENTRY-2502](#) - Modified BackportSentry NN plug-in stops fetching updates from sentry server.
- [SENTRY-2511](#) - Debug level logging on HMSPaths significantly affects performance

Apache Spark

- [SPARK-4224](#) - [CORE][YARN] Support group acls
- [SPARK-19019](#) - [PYTHON][BRANCH-1.6] Fix hijacked `collections.namedtuple` and port cloudpickle changes for PySpark to work with Python 3.6.0

Apache Zookeeper

- [ZOOKEEPER-1392](#) - Request READ or ADMIN permission for getAcl()
- [ZOOKEEPER-2141](#) - ACL cache in DataTree never removes entries

Issues Fixed in CDH 5.16.1

CDH 5.16.1 fixes the following issues:

CVE-2019-10099: Apache Spark local files left unencrypted

Certain operations in Spark leave local files unencrypted on disk, even when local file encryption is enabled with “spark.io.encryption.enabled”.

This includes cached blocks that are fetched to disk (controlled by `spark.maxRemoteBlockSizeFetchToMem`) in the following cases:

- In SparkR when parallelize is used
- In Pyspark when broadcast and parallelize are used
- In Pyspark when python udfs is used

Products affected:

- CDH
- CDS Powered by Apache Spark

Affected versions:

- CDH 5.15.1 and earlier
- CDH 6.0.0
- CDS 2.1.0 release 1 and release 2
- CDS 2.2.0 release 1 and release 2
- CDS 2.3.0 release 3

Users affected: All users who run Spark on CDH and CDS in a multi-user environment.

Date/time of detection: July 2018

Severity (Low/Medium/High): 6.3 Medium ([CVSS AV:L/AC:H/PR:N/UI:R/S:U/C:H/I:H/A:N](#))

Impact: Unencrypted data accessible.

CVE: CVE-2019-10099

Immediate action required: Upgrade to a version of CDH containing the fix.

Workaround: Do not use of pyspark and the fetch-to-disk options.

Fixed versions:

- CDH 5.15.2
- CDH 5.16.0
- CDH 6.0.1
- CDS 2.1.0 release 3
- CDS 2.2.0 release 3
- CDS 2.3.0 release 4

For the latest update on this issue see the corresponding Knowledge article: [TSB 20210-336: Apache Spark local files left unencrypted](#)

[CVE-2018-1296 Permissive Apache Hadoop HDFS listXAttr Authorization Exposes Extended Attribute Key/Value Pairs](#)

AHDFS exposes extended attribute key/value pairs during listXAttrs, verifying only path-level search access to the directory rather than path-level read permission to the referent.

Products affected: Apache HDFS

Releases affected:

- CDH 5.4.0 - 5.15.1, 5.16.0
- CDH 6.0.0, 6.0.1, 6.1.0

Users affected: Users who store sensitive data in extended attributes, such as users of HDFS encryption.

Date/time of detection: December 12, 2017

Detected by: Rushabh Shah, Yahoo! Inc., Hadoop committer

Severity (Low/Medium/High): Medium

Impact: HDFS exposes extended attribute key/value pairs during listXAttrs, verifying only path-level search access to the directory rather than path-level read permission to the referent. This affects features that store sensitive data in extended attributes.

CVE: CVE-2018-1296

Immediate action required:

- Upgrade: Update to a version of CDH containing the fix.
- Workaround: If a file contains sensitive data in extended attributes, users and admins need to change the permission to prevent others from listing the directory that contains the file.

Addressed in release/refresh/patch:

- CDH 5.15.2, 5.16.1
- CDH 6.1.1, 6.2.0

[Apache Tomcat Vulnerability CVE-2018-11784](#)

Fixed a vulnerability in Apache Tomcat where specially-crafted URLs could be used to redirect to any given URI. CVE-2018-11784.

Cloudera Issue: CDH-73885

[Cloudera Search restore operation puts shard replicas on same host](#)

Restoring an Apache Solr collection sometimes places all shard replicas on the same host.

Cloudera Issue: CDH-68828

[Missing authorization in Apache Impala may allow data injection](#)

A malicious user who is authenticated with Kerberos may have unauthorized access to internal services used by Impala to transfer intermediate data during query execution. If details of a running query (e.g. query ID, query plan) are available, a user can craft some RPC requests with custom software to inject data into a running query or end query execution prematurely, leading to wrong results of the query.

Cloudera Issue: CDH-72373 / TSB-338

CVE: CVE-2018-11785

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.16.1:

Apache Hadoop

- [HADOOP-13426](#) - Improved IPC performance.
- [HADOOP-13483](#) - Fixed an issue where file-create overwrote directories instead of throwing error messages
- [HADOOP-15473](#) - Configured the serialFilter in KeyProvider to avoid UnrecoverableKeyException caused by JDK-8189997
- [HADOOP-15655](#) - Enhanced KMS client retry behavior. Previously, the KMS did not retry upon SocketTimeoutException.
- [HDFS-8229](#) - Fixed an issue where the LAZY_PERSIST file gets deleted after NameNode restart.
- [HDFS-10240](#) - Fixed a race between close/recoverLease leads to missing block
- [HDFS-12299](#) - Fixed a race between update pipeline and DN Re-Registration
- [HDFS-13051](#) - Fixed a dead lock during async editlog rolling if the edit queue is full.
- [HDFS-13322](#) - Fixed an issue where the UID persists when switching between ticket caches.
- [HDFS-13486](#) - Fixed an issue where a faulty node can cause a lease leak and NPE on accessing data.
- [HDFS-13601](#) - Optimized ByteString conversions in PBHelper.
- [HDFS-13611](#) - Fixed an issue where text was used as a ConcurrentHashMap key in PBHelperClient.
- [HDFS-13813](#) - Added a check to see if a child inode exists in the global FSDirectory directory when saving (serializing) INodeDirectorySection.
- [MAPREDUCE-7053](#) - Fixed an issue where Timed out tasks can fail to produce thread dump
- [YARN-6966](#) - Fixed an issue where NodeManager metrics may return wrong negative values when NM restart.
- [YARN-6967](#) - Fixed an issue where the limit for diagnostic message size was not honored
- [YARN-8436](#) - Fixed an issue where the ResourceManager can fail while sorting queues if an update comes in

Apache HBase

- [HBASE-15232](#) - Handle region location cache mgmt in AsyncProcess for multi()'s
- [HBASE-15390](#) - Unnecessary MetaCache evictions cause elevated number of requests to meta
- [HBASE-18891](#) - Upgrade to netty-all 4.0.50.Final
- [HBASE-19924](#) - hbase rpc throttling does not work for multi() with request count rater.
- [HBASE-20493](#) - Port HBASE-19994 (Create a new class for RPC throttling exception, make it retryable) to branch-1
- [HBASE-20651](#) - Master, prevents hbck or shell command to reassign the split parent region
- [HBASE-20723](#) - Custom hbase.wal.dir results in data loss because we write recovered edits into a different place than where the recovering region server looks for them
- [HBASE-20997](#) - rebuildUserRegions() does not build ReplicaMapping during master switchover

Apache Hive

Code Changes Should Not Be Required

The following fixes should not require code changes, but they contain improvements that might enhance your deployment:

- [HIVE-6980](#) - Drop table by using direct sql
- [HIVE-10296](#) - Cast exception observed when hive runs a multi join query on metastore (postgres), since postgres pushes the filter into the join, and ignores the condition before applying cast
- [HIVE-12981](#) - ThriftCLIService uses incompatible getShortName() implementation
- [HIVE-15237](#) - Propagate Spark job failure to Hive
- [HIVE-15860](#) - RemoteSparkJobMonitor may hang when RemoteDriver exits abnormally
- [HIVE-16483](#) - HoS should populate split related configurations to HiveConf
- [HIVE-17213](#) - HoS file merging doesn't work for union all
- [HIVE-18031](#) - Support replication for Alter Database operation
- [HIVE-18283](#) - Better error message and error code for HoS exceptions

- [HIVE-18765](#) - SparkClientImpl swallows exception messages from the RemoteDriver
- [HIVE-18916](#) - SparkClientImpl doesn't error out if spark-submit fails
- [HIVE-19259](#) - Create view on tables having union all fail with 'Table not found'
- [HIVE-19310](#) - Metastore: MetaStoreDirectSql.ensureDbInit has some slow DN calls which might need to be run only in test env
- [HIVE-19371](#) - Add table ownerType to HMS thrift API
- [HIVE-19372](#) - Add table ownerType to JDO/SQL and ObjectStore
- [HIVE-19374](#) - Parse and process ALTER TABLE SET OWNER command syntax
- [HIVE-19605](#) - TAB_COL_STATS table has no index on db/table name
- [HIVE-19668](#) - Over 30% of the heap wasted by duplicate org.antlr.runtime.CommonToken's and duplicate strings
- [HIVE-19783](#) - Retrieve only locations in HiveMetaStore.dropPartitionsAndGetLocations
- [HIVE-20183](#) - Inserting from bucketed table can cause data loss, if the source table contains empty bucket
- [HIVE-20345](#) - Drop database may hang if the tables get deleted from a different call

Hue

- [HUE-8118](#) - [core] Fine grain tracking of the memory usage
- [HUE-8118](#) - [core] The duration of the request is always shown even when instrumentation flag is off
- [HUE-8128](#) - [backend] Force debug logging in server logs does not get all debug
- [HUE-8162](#) - [core] Add delete operation to the right document assist
- [HUE-8177](#) - [oozie] Add a config check for /user/hue/oozie/workspaces
- [HUE-8377](#) - [security] Support new Sentry finer grain privileges
- [HUE-8377](#) - [security] Correctly apply the new permissions to the database scope
- [HUE-8451](#) - [notebook] Many "codec can't decode byte" errors on pig execution if browser language=jp
- [HUE-8464](#) - [core] Fix SAML encryption missing key file passphrase
- [HUE-8467](#) - [jobbrowser] Support impala digest auth for queries
- [HUE-8475](#) - [report] Protect against pivot conflicting with nested facets
- [HUE-8476](#) - [frontend] Fix jQuery Hive autocomplete column mapping
- [HUE-8487](#) - [useradmin] Fix Add Sync LDAP user fails when using DN with special character
- [HUE-8505](#) - [core] Close impala session on logout
- [HUE-8519](#) - [jb] Impala API can now directly return json
- [HUE-8558](#) - [jb] Add tracking URL to Spark Jobs and remove url and killUrl
- [HUE-8564](#) - [useradmin] Fix last activity update for jobbrowser/api/jobs requests
- [HUE-8564](#) - [useradmin] Fix last activity update for notebook/api/check_status
- [HUE-8571](#) - [sentry] navigator_api ERROR for PRIVILEGE_HIERARCHY[hierarchy[server][SENTRY_PRIVILEGE_KEY]['action']]
- [HUE-8602](#) - [sentry] Remove ALTER and DROP in the Hive section

Apache Impala

- [IMPALA-6086](#) - Require the SELECT privilege on the database for built-in function calls.
- [IMPALA-6451](#) - Fixed the AuthorizationException in CTAS for Kudu tables.
- [IMPALA-6479](#) - DESCRIBE now respects column level privileges and only shows the columns that the user has the privilege to view.
- [IMPALA-6571](#) - Fixed the NullPointerException in SHOW CREATE TABLE for HBase tables.
- [IMPALA-7225](#) - REFRESH..PARTITION no longer reset the number of rows in a partition.
- [IMPALA-7272](#) - Fixed the crash in StringMinMaxFilter.
- [IMPALA-7360](#) - Fixed an issue where Avro scanner sometimes skipped blocks when skip marker was on HDFS block boundary.
- [IMPALA-7419](#) - Fixed the NullPointerException in SimplifyConditionalsRule.
- [IMPALA-7483](#) - impalad/catalogd on JVM deadlock now get aborted.
- [IMPALA-7520](#) - Fixed the NullPointerException in SentryProxy.

Apache Kudu

- [KUDU-2260](#) - Fixed a rare issue where system failure could leave unexpected null bytes at the end of metadata files, causing Kudu to be unable to restart.
- [KUDU-2364](#) - Fixed an issue when a tablet server was wiped and recreated with the same RPC address, ksck listed it twice, both as healthy, even though only one of them was there.
- [KUDU-2412](#) - The kudu-python client can now compile in environments where `__int128` is not supported. This was most commonly el6 environments.
- [KUDU-2509](#) - Fixed an issue that might result in a crash of a tablet server in case of a WAL replay error while bootstrapping a tablet.
- [KUDU-2580](#) - Fixed authentication token reacquisition in the C++ client.
- Fixed an issue that caused the kudu CLI tool to unexpectedly exit when the connection to the master or tserver was abruptly closed.

Apache Oozie

- [OOZIE-2457](#) - Oozie log parsing regex consume more than 90% cpu
- [OOZIE-3193](#) - Applications are not killed when submitted via subworkflow
- [OOZIE-3354](#) - [core] [SSH action] SSH action gets hung
- [OOZIE-3370](#) - Property filtering is not consistent across job submission

Apache Pig

The following issues are fixed in CDH 6.1.1:

- [PIG-5373](#) - InterRecordReader might skip records if certain sync markers are used
- [PIG-5374](#) - Use CircularFifoBuffer in InterRecordReader

Apache Sentry

- [SENTRY-1272](#) - Enable ALTERVIEW_RENAME and ALTERVIEW_AS operation in hive binding
- [SENTRY-2194](#) - Upgrade Sentry hadoop-version dependency to 2.7.5
- [SENTRY-2210](#) - AUTHZ_PATH should have index on the foreign key AUTHZ_OBJ_ID
- [SENTRY-2214](#) - Sentry should not allow URI grants to EMPTY or NULL locations
- [SENTRY-2219](#) - Create index AUTHZ_PATH_FK_IDX at table AUTHZ_PATH only when it does not exist for Oracle
- [SENTRY-2238](#) - Explicitly set Database on SentryHivePrivilegeObjectDesc
- [SENTRY-2299](#) - NPE In Sentry HDFS Sync Plugin
- [SENTRY-2310](#) - Sentry is not be able to fetch full update subsequently, when there is HMS restart in the snapshot process.
- [SENTRY-2332](#) - Load hadoop default configuration when starting sentry service
- [SENTRY-2333](#) - Create index AUTHZ_PATH_FK_IDX at table AUTHZ_PATH for Postgres only when it does not exist
- [SENTRY-2403](#) - Incorrect naming in RollingFileWithoutDeleteAppender
- [SENTRY-2406](#) - Make sure inputHierarchy and outputHierarchy have unique values

Apache Solr

- [SOLR-12290](#) - Do not close any servlet streams and improve our servlet stream closing prevention code for users and devs.
- [SOLR-12293](#) - Updates need to use their own connection pool to maintain connection reuse and prevent spurious recoveries.

Apache Spark

- [SPARK-22864](#) - [CORE] Disable allocation schedule in ExecutorAllocationManagerSuite.
- [SPARK-25253](#) - [PYSPARK] Refactor local connection & auth code
- [SPARK-25318](#) - Add exception handling when wrapping the input stream during the the fetch or stage retry in response to a corrupted block

Apache Zookeeper

- [ZOOKEEPER-706](#) - Large numbers of watches can cause session re-establishment to fail
- [ZOOKEEPER-1382](#) - Zookeeper server holds onto dead/expired session ids in the watch data structures

Issues Fixed in CDH 5.15.x

The following topics describe issues fixed in CDH 5.15.x, from newest to oldest release. You can also review [What's New in CDH 5.15.x](#) on page 1814 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.15.2

CVE-2019-10099: Apache Spark local files left unencrypted

Certain operations in Spark leave local files unencrypted on disk, even when local file encryption is enabled with “spark.io.encryption.enabled”.

This includes cached blocks that are fetched to disk (controlled by `spark.maxRemoteBlockSizeFetchToMem`) in the following cases:

- In SparkR when parallelize is used
- In Pyspark when broadcast and parallelize are used
- In Pyspark when python udfs is used

Products affected:

- CDH
- CDS Powered by Apache Spark

Affected versions:

- CDH 5.15.1 and earlier
- CDH 6.0.0
- CDS 2.1.0 release 1 and release 2
- CDS 2.2.0 release 1 and release 2
- CDS 2.3.0 release 3

Users affected: All users who run Spark on CDH and CDS in a multi-user environment.

Date/time of detection: July 2018

Severity (Low/Medium/High): 6.3 Medium ([CVSS AV:L/AC:H/PR:N/UI:R/S:U/C:H/I:H/A:N](#))

Impact: Unencrypted data accessible.

CVE: CVE-2019-10099

Immediate action required: Upgrade to a version of CDH containing the fix.

Workaround: Do not use of pyspark and the fetch-to-disk options.

Fixed versions:

- CDH 5.15.2
- CDH 5.16.0
- CDH 6.0.1
- CDS 2.1.0 release 3
- CDS 2.2.0 release 3
- CDS 2.3.0 release 4

For the latest update on this issue see the corresponding Knowledge article: [TSB 20210-336: Apache Spark local files left unencrypted](#)

[CVE-2018-1296 Permissive Apache Hadoop HDFS listXAttr Authorization Exposes Extended Attribute Key/Value Pairs](#)

AHDFS exposes extended attribute key/value pairs during listXAttrs, verifying only path-level search access to the directory rather than path-level read permission to the referent.

Products affected: Apache HDFS

Releases affected:

- CDH 5.4.0 - 5.15.1, 5.16.0
- CDH 6.0.0, 6.0.1, 6.1.0

Users affected: Users who store sensitive data in extended attributes, such as users of HDFS encryption.

Date/time of detection: December 12, 2017

Detected by: Rushabh Shah, Yahoo! Inc., Hadoop committer

Severity (Low/Medium/High): Medium

Impact: HDFS exposes extended attribute key/value pairs during listXAttrs, verifying only path-level search access to the directory rather than path-level read permission to the referent. This affects features that store sensitive data in extended attributes.

CVE: CVE-2018-1296

Immediate action required:

- Upgrade: Update to a version of CDH containing the fix.
- Workaround: If a file contains sensitive data in extended attributes, users and admins need to change the permission to prevent others from listing the directory that contains the file.

Addressed in release/refresh/patch:

- CDH 5.15.2, 5.16.1
- CDH 6.1.1, 6.2.0

[Missing authorization in Apache Impala may allow data injection](#)

A malicious user who is authenticated with Kerberos may have unauthorized access to internal services used by Impala to transfer intermediate data during query execution. If details of a running query (e.g. query ID, query plan) are available, a user can craft some RPC requests with custom software to inject data into a running query or end query execution prematurely, leading to wrong results of the query.

Cloudera Issue: CDH-72373 / TSB-338

CVE: CVE-2018-11785

[Upstream Issues Fixed](#)

The following upstream issues are fixed in CDH 5.15.2:

Apache Hadoop

- [HADOOP-15473](#) - Enhanced serialFilter in KeyProvider to avoid UnrecoverableKeyException caused by JDK-8189997.
- [HADOOP-15655](#) - Enhanced KMS client retry behavior.
- [HDFS-10240](#) - Fixed an issue where a race between close/recoverLease leads to missing blocks.
- [HDFS-12299](#) - Fixed a race between update pipeline and DataNode Re-Registration.
- [HDFS-13322](#) - Fixed an issue with fuse dfs that caused the UID to persist when switching between ticket caches.
- [HDFS-13486](#) - Backport HDFS-11817 to fix an issue where a faulty node can cause a lease leak and NPE on accessing data.
- [HDFS-13813](#) - Enhanced NameNode behavior to exit if dangling child inode is detected when saving FsImage.
- [MAPREDUCE-7053](#) - Fixed an issue where timed out tasks can fail to produce thread dump.
- [YARN-6966](#) - Fixed an issue where NodeManager metrics may return wrong negative values when NM restart.
- [YARN-6967](#) - Fixed an issue where an application attempt's diagnostic message size was not properly limited.

- [YARN-8436](#) - Fixed an issue with FSParentQueue where the comparison method violates its general contract.

Apache HBase

- [HBASE-19730](#) - HBASE-14497 Reverse Scan threw StackOverflow caused by readPt checking
- [HBASE-19924](#) - hbase rpc throttling does not work for multi() with request count rater.
- [HBASE-20493](#) - Port HBASE-19994 (Create a new class for RPC throttling exception, make it retryable) to branch-1
- [HBASE-20723](#) - Custom hbase.wal.dir results in data loss because we write recovered edits into a different place than where the recovering region server looks for them.
- [HBASE-20997](#) - rebuildUserRegions() does not build ReplicaMapping during master switchover
- [HBASE-21275](#) - Disable TRACE HTTP method for thrift http server

Apache Hive

- [HIVE-12981](#) - ThriftCLIService uses incompatible getShortName() implementation
- [HIVE-13394](#) - Analyze table fails in tez on empty partitions
- [HIVE-14236](#) - Partial backport of to fix union-stats related errors
- [HIVE-14560](#) - Support exchange partition between s3 and hdfs tables
- [HIVE-14690](#) - Query fail when hive.exec.parallel=true, with conflicting session dir
- [HIVE-16483](#) - HoS should populate split related configurations to HiveConf
- [HIVE-17213](#) - HoS: file merging doesn't work for union all
- [HIVE-19259](#) - Create view on tables having union all fail with 'Table not found'
- [HIVE-20183](#) - Inserting from bucketed table can cause data loss, if the source table contains empty bucket
- [HIVE-20345](#) - Drop database may hang if the tables get deleted from a different call
- [HIVE-20678](#) - HiveHBaseTableOutputFormat should implement HiveOutputFormat to ensure compatibility
- [HIVE-20695](#) - HoS Query fails with hive.exec.parallel=true

Hue

- [HUE-8128](#) - [backend] Force debug logging in server logs does not get all debug
- [HUE-8398](#) - [editor] Fix broken result table after multiple queries in embedded mode
- [HUE-8399](#) - [editor] Various improvements for embedded mode
- [HUE-8399](#) - [editor] Fix various embedded mode issues in old version
- [HUE-8399](#) - [editor] Limit functions assist to Impala in embedded mode
- [HUE-8399](#) - [editor] Hide the right assistant in embedded mode
- [HUE-8451](#) - [notebook] Many "codec can't decode byte" errors on pig execution if browser language=jp
- [HUE-8458](#) - [frontend] Fix issue with async loading of js resources in the dashboard
- [HUE-8458](#) - [frontend] Evaluate the js resources while others are being fetched
- [HUE-8458](#) - [frontend] Load new scripts using \$.get and eval instead of appending <script> tags
- [HUE-8464](#) - [core] Fix SAML encryption missing key file passphrase
- [HUE-8467](#) - [jobbrowser] Support impala digest auth for queries
- [HUE-8468](#) - [frontend] Append a style tag to head instead of modifying stylesheets for dynamic styles in embedded mode
- [HUE-8475](#) - [report] Protect against pivot conflicting with nested facets
- [HUE-8487](#) - [useradmin] Fix Add Sync LDAP user fails when using DN with special character
- [HUE-8505](#) - [core] Close impala session on logout
- [HUE-8558](#) - [jb] Add tracking URL to Spark Jobs and remove url and killUrl
- [HUE-8564](#) - [useradmin] Fix last activity update for notebook/api/check_status
- [HUE-8564](#) - [useradmin] Fix last activity update for jobbrowser/api/jobs requests

Apache Impala

- [IMPALA-6907](#) - Now Impala correctly closes all stale connections to removed impala cluster members.
- [IMPALA-7225](#) - Fixed an issue where the REFRESH...PARTITION statement caused statistics for the refreshed partition to be automatically reset to -1 (unknown) . With the fix, statistics will be changed only if an explicit COMPUTE STATS statement is issued for an object.

- [IMPALA-7272](#) - Fixed a crash caused by a memory management problem when the query execution requires finding strings inside a range defined by the lesser-than and greater-than comparisons.
- [IMPALA-7360](#) - Fixed an issue where Impala could incorrectly skip data if a record separator in a sequence-based file (Avro, RC or sequence file) straddled an HDFS block boundary.
- [IMPALA-7537](#) - Fixed a security issue where REVOKE ALL ON SERVER did not have a permanent effect if the ALL permission was granted using the WITH GRANT option. Running INVALIDATE METADATA no longer causes the permission to reappear.
- [IMPALA-7585](#) - Fixed an issue in KRPC, which can cause slow or hung queries on non-secured clusters.

Apache Kudu

- [KUDU-2463](#) - Fixed an issue in which incorrect results would be returned in scans following a server restart.
- [KUDU-2509](#) - Fixed an issue that might result in a crash of a tablet server in case of a WAL replay error while bootstrapping a tablet.
- [KUDU-2580](#) - Fixed authentication token reacquisition in the C++ client.

Apache Oozie

- [OOZIE-2457](#) - Oozie log parsing regex consume more than 90% cpu
- [OOZIE-3354](#) - [core] [SSH action] SSH action gets hung
- [OOZIE-3370](#) - Property filtering is not consistent across job submission

Apache Sentry

- [SENTRY-1944](#) - Optimize DelegateSentryStore.getGroupsByRoles() and update SentryGenericPolicyProcessor to retrieve roles to group mapping in a single transaction
- [SENTRY-2194](#) - Upgrade Sentry hadoop-version dependency to 2.7.5
- [SENTRY-2214](#) - Sentry should not allow URI grants to EMPTY or NULL locations
- [SENTRY-2309](#) - ModifiedCatch NPE thrown when fetching Partitions with no corresponding SDS entry
- [SENTRY-2332](#) - Load hadoop default configuration when starting sentry service
- [SENTRY-2403](#) - Incorrect naming in RollingFileWithoutDeleteAppender
- [SENTRY-2406](#) - Make sure inputHierarchy and outputHierarchy have unique values
- [SENTRY-2428](#) - Skip null partitions or partitions with null sds entries

Apache Spark

- [SPARK-25253](#) - [PYSPARK] Refactor local connection & auth code
- [SPARK-25318](#) - Add exception handling when wrapping the input stream during the the fetch or stage retry in response to a corrupted block

Apache Zookeeper

- [ZOOKEEPER-706](#) - Large numbers of watches can cause session re-establishment to fail
- [ZOOKEEPER-1382](#) - Zookeeper server holds onto dead/expired session ids in the watch data structures

Issues Fixed in CDH 5.15.1

CDH 5.15.1 fixes the following issues.

Hadoop YARN Privilege Escalation CVE-2016-6811

Fixed a vulnerability in Hadoop YARN that allows a user who can escalate to the yarn user the ability to possibly run arbitrary commands as the root user. CVE-2016-6811

Apache HBase potential privilege escalation for user of HBase "Thrift 1" API Server over HTTP CVE-2018-8025

CVE-2018-8025 describes an issue in Apache HBase that affects the optional "Thrift 1" API server when running over HTTP. There is a race-condition that could lead to authenticated sessions being incorrectly applied to users, e.g. one authenticated user would be considered a different user or an unauthenticated user would be treated as an authenticated user.

Products affected: HBase Thrift Server

Releases affected:

- CDH 5.4.x - 5.12.x
- CDH 5.13.0, 5.13.1, 5.13.2, 5.13.3
- CDH 5.14.0, 5.14.2, 5.14.3
- CDH 5.15.0

Fixed versions: CDH 5.14.4, 5.15.1

Users affected: Users with the HBase Thrift 1 service role installed and configured to work in “thrift over HTTP” mode. For example, those using Hue with HBase impersonation enabled.

Severity: High

Potential privilege escalation.

CVE: CVE-2018-8025

Immediate action required: Upgrade to a CDH version with the fix, or, disable the HBase Thrift-over-HTTP service. Disabling the HBase Thrift-over-HTTP service will render Hue impersonation inoperable and all HBase access via Hue will be performed using the “hue” user instead of the authenticated user.

Knowledge article: For the latest update on this issue see the corresponding Knowledge article - TSB: 2018-315: Potential privilege escalation for user of HBase “Thrift 1” API Server over HTTP

[Apache Hive CREATE TABLE EXTERNAL commands run very slowly on Amazon S3](#)

When you run a `CREATE EXTERNAL TABLE` command on Amazon S3 storage, the command runs extremely slowly and might not complete. This usually occurs for tables that have pre-existing, nested data. There is no known workaround.

Affected Versions: CDH 5.15.0

Cloudera Bug: CDH-68833

Workaround: None

Fixed versions: CDH 5.15.1 and later

[Cloudera Search restore operation puts shard replicas on same host](#)

Restoring an Apache Solr collection sometimes places all shard replicas on the same host.

Cloudera Issue: CDH-68828

[Zip Slip Vulnerability CVE-2018-8009](#)

“Zip Slip” is a widespread arbitrary file overwrite critical vulnerability, which typically results in remote command execution. It was discovered and responsibly disclosed by the Snyk Security team ahead of a public disclosure on June 5, 2018, and affects thousands of projects.

Cloudera has analyzed our use of zip-related software, and has determined that only Apache Hadoop is vulnerable to this class of vulnerability in CDH 5. This has been fixed in upstream Hadoop as CVE-2018-8009.

[hdfs snapshotDiff /.reserved/raw/... fails on snapshottable directories](#)

Fixed an issue where the `hdfs snapshotDiff` command fails when the command is supplied with the raw path.

Cloudera Issue: CDH-66029

[Upstream Issues Fixed](#)

The following upstream issues are fixed in CDH 5.15.1:

Apache Flume

- [FLUME-2786](#), [FLUME-3056](#), [FLUME-3117](#) - Application enters a deadlock when stopped while `handleConfigurationEvent`
- [FLUME-2894](#) - Flume components should stop in the correct order
- [FLUME-2976](#) - Exception when JMS source tries to connect to a Weblogic server without authentication

- [FLUME-3222](#) - Fix for NoSuchFileException thrown when files are being deleted

Apache Hadoop

- [HADOOP-13024](#) - Distcp with -delete feature on raw data not implemented
- [HADOOP-13972](#) - ADLS to support per-store configuration
- [HADOOP-15186](#) - Allow Azure Data Lake SDK dependency version to be set on the command line
- [HADOOP-15317](#) - Improve NetworkTopology chooseRandom's loop
- [HADOOP-15342](#) - Updating ADLS connector to use SDK version 2.2.7
- [HADOOP-15356](#) - Make the HTTP timeout configurable in ADLS connector
- [HADOOP-15434](#) - Upgrade to ADLS SDK that exposes current timeout
- [HADOOP-15466](#) - Correct units in `adl.http.timeout` property
- [HDFS-9229](#) - Expose size of NameNode directory as a metric
- [HDFS-11751](#) - DFSZKFailoverController daemon exits with wrong status code
- [HDFS-11993](#) - Add log info when connect to datanode socket address failed
- [HDFS-12683](#) - DFSZKFailOverController re-order logic for logging Exception
- [HDFS-12710](#) - HTTPFS HTTP max header size env variable is not respected in branch-2
- [HDFS-12981](#) - Running `renameSnapshot` on a non-existent snapshot to itself should throw error
- [HDFS-13281](#) - `Namenode#createFile` should be `/.reserved/raw/` aware
- [HDFS-13314](#) - NameNode should optionally exit if it detects Fslmage corruption
- [MAPREDUCE-7094](#) - LocalDistributedCacheManager leaves classloaders open, which leaks File Descriptors
- [YARN-4227](#) - Ignore expired containers from removed nodes in FairScheduler
- [YARN-4325](#) - NodeManager log handlers fail to send finished/failed events in some cases
- [YARN-4677](#) - `RMNodeResourceUpdateEvent` update from scheduler can lead to a race condition
- [YARN-5121](#) - Fix container-executor portability issues

Apache HBase

Code Changes Should Not Be Required

The following fixes should not require code changes, but they contain improvements that might enhance your deployment:

- [HBASE-20229](#) - `ConnectionImplementation.locateRegions()` returns duplicated entries when region replication is on
- [HBASE-20293](#) - `get_splits` returns duplicate split points when region replication is on
- [HBASE-20664](#) - Reduce the broad scope of `outToken` in `ThriftHttpServlet`

Apache Hive

Code Changes Should Not Be Required

The following fixes should not require code changes, but they contain improvements that might enhance your deployment:

- [HIVE-9915](#) - Allow specifying file format for managed tables
- [HIVE-15580](#) - Eliminate unbounded memory usage for `orderBy` and `groupBy` in Hive on Spark
- [HIVE-15682](#) - Eliminate per-row based dummy iterator creation
- [HIVE-15683](#) - Make configurable how Hive on Spark shuffles data for `GROUP BY`
- [HIVE-16080](#) - Add `parquet` to possible values for `hive.default.fileformat` and `hive.default.fileformat.managed`
- [HIVE-16659](#) - Query plan should reflect `hive.spark.use.groupby.shuffle`
- [HIVE-19041](#) - Thrift deserialization of `Partition` objects should intern fields
- [HIVE-19231](#) - Beeline generates garbled output when using `UnsupportedTerminal`
- [HIVE-19424](#) - NPE In `MetaDataFormatters`
- [HIVE-19668](#) - Over 30% of the heap is wasted by duplicate `org.antlr.runtime.CommonTokens` and duplicate strings
- [HIVE-19700](#) - Workaround for `JLine` issue with `UnsupportedTerminal`
- [HIVE-19870](#) - `HCatalog` dynamic partition query can fail if the table path is managed by Sentry

Hue

- [HUE-6697](#) - [jb] Prevent reset of job page tabs when job is running
- [HUE-7946](#) - [jb] Link to subworkflow on workflow dashboard page 404s in Hue 4
- [HUE-7989](#) - [useradmin] Provide better UI message and message in logs when ldap server down
- [HUE-8053](#) - Fix unit test
- [HUE-8053](#) - [useradmin] LDAP authentication with sync_groups_on_login=true fails with KeyError exception
- [HUE-8055](#) - [desktop] Support multiple LDAP servers in LDAP Test command
- [HUE-8063](#) - [fb] Better error message when HTTPFS is not working
- [HUE-8137](#) - [core] Refresh translations files for 4.2
- [HUE-8159](#) - [oozie] Unable to create Workflow/Schedule using Java document Action
- [HUE-8170](#) - [useradmin] Fix LDAP sync (ldap_access.py) certificate validation logic
- [HUE-8173](#) - [core] Add a warning log when a users switched to the old Hue 3 UI
- [HUE-8173](#) - [core] Add a log line when a user loads a page not via the load balancer even if we have one
- [HUE-8191](#) - [jb] UnicodeEncodeError occurs in Job Browser when browser language is not English
- [HUE-8202](#) - [jb] Fix mutual authentication failed with Isilon
- [HUE-8207](#) - [indexer] Issue previewing input file with unicode data
- [HUE-8217](#) - [dashboard] Fix HTML resultset widget templating
- [HUE-8232](#) - [oozie] Fix 500 error if coordinator associated workflow has been deleted
- [HUE-8236](#) - [core] Correct Hue config value use integer instead of math expression
- [HUE-8253](#) - [editor] Support downloading Query results with query names(file names) other than ISO-8859-1 charset
- [HUE-8253](#) - [editor] Fix the broken unit test in 5.15 due to conflict after backporting
- [HUE-8274](#) - [s3] Moving a folder using drag and drop deletes the folder itself
- [HUE-8280](#) - [fb] Move action button does not prevent move to itself
- [HUE-8280](#) - [fb] Fix the missed conflict issue after backporting
- [HUE-8282](#) - [fb] Check if EC2 instance before check IAM metadata
- [HUE-8305](#) - [useradmin] Optimize performance on checking Hue permissions if user is in many groups
- [HUE-8310](#) - Broken template for multiple custom apps
- [HUE-8313](#) - [core] Remove hardcoding to ImpersonationBackend when using embedded mode
- [HUE-8314](#) - [core] Fix SAML encryption missing config
- [HUE-8320](#) - [core] After copying a workflow using Hue 4 button, saving the copied workflow fails
- [HUE-8344](#) - [hbase] Hbase old version of data can not display in Hue
- [HUE-8350](#) - [solr] indexer app permission is not being acknowledged in HUE
- [HUE-8350](#) - [solr] indexer app permission is not being acknowledged in HUE
- [HUE-8370](#) - [pig] Imported old version Pig script missing properties fields
- [HUE-8392](#) - [oozie] Cannot add more actions using drag & drop from actions bar in the Oozie editor after adding around 3 actions
- [HUE-8404](#) - [useradmin] Fix multibackend invalid password removes drop down to select Local
- [HUE-8409](#) - [core] When idle session timeout is enabled it causes issues with Spnego
- [HUE-8440](#) - [jb] Link for Spark logs in Properties tab of Job Browser is incorrect
- [HUE-8455](#) - [pig] Oozie editor fails with 'hadoopProperties' for pig script saved in Hue 3

Apache Impala

- [IMPALA-6687](#) - Fixed INSERT with mixed case in partition column names.
- [IMPALA-6822](#) - Added a query option to control shuffling by distinct expressions.
- [IMPALA-6847](#) - Added a work around high memory estimates for admission control.
- [IMPALA-6908](#) - IsConnResetTException() should include ECONNRESET.
- [IMPALA-6934](#) - Corrected the wrong results with EXISTS subqueries that contain ORDER BY, LIMIT, and OFFSET.
- [IMPALA-7014](#) - Disabled the stacktrace symbolisation by default.
- [IMPALA-7078](#) - Reduced the queue size based on num_scanner_threads.

- [IMPALA-7078](#) - Improved memory consumption of wide Avro scans.
- [IMPALA-7288](#) - Fixed the Codegen crash in FinalizeModule.
- [IMPALA-7298](#) - Impala no longer passes IP address as hostname in Kerberos principal.

Apache Kudu

- [KUDU-2367](#) - Fixed an issue where a permanently failed tablet replica was not properly identified, which could cause the tablet not to re-replicate in very small clusters.
- [KUDU-2377](#) - Fixed an issue that caused Kudu servers to fail to start when RLIMIT_NPROC=-1.
- [KUDU-2378](#) - Fixed unaligned loads of int128 from rows.
- [KUDU-2379](#) - Fixed an issue that caused secure Spark jobs to fail.
- [KUDU-2416](#) - Fixed PartialRow.setMin.
- [KUDU-2443](#) - Fixed replica movement and replacement for RF=1.
- [KUDU-2447](#) - Fixed the tablet server crash with the error, "NONE predicate can not be pushed into key".
- [KUDU-2478](#) - Restored Python 2.6 compatibility.
- Added the ability to adjust scan timeouts in Spark.
- Increased the timeout to begin tablet copies, which improves Kudu's re-replication time when the cluster is busy.
- Fixed a NullPointerException thrown when calling ColumnSchema#toString on non-decimal types.
- Greatly improved the performance of many types of queries on tables from which many rows have been deleted.
- Fixed an issue that caused partition pruning to be too conservative for queries from the Java client that use Decimal predicates.

Apache Oozie

- [OOZIE-2491](#) - oozie acl cannot specify group, it does not work
- [OOZIE-3134](#) - Potential inconsistency between the in-memory SLA map and the Oozie database
- [OOZIE-3260](#) - [sla] Remove stale item above max retries on JPA related errors from in-memory SLA map

Apache Parquet

- [PARQUET-1246](#) - Ignore float/double statistics in case of NaN

Apache Sentry

Code Changes Should Not Be Required

The following fixes should not require code changes, but they contain improvements that might enhance your deployment:

- [SENTRY-1209](#) - Sentry does not block Hive's cross-schema table renames
- [SENTRY-2020](#) - Fix testConsumeCycleWithInsufficientPrivileges test failure in kafka e2e tests
- [SENTRY-2144](#) - Table Rename Cross Database should update permission correctly
- [SENTRY-2165](#) - NotificationProcessor process notification methods have logs wrongly flagged as ERROR
- [SENTRY-2183](#) - Increase default sentry-hdfs rpc timeout to 20 mins
- [SENTRY-2184](#) - Performance Issue: MPath is queried for each MAuthzPathsMapping in full snapshot
- [SENTRY-2226](#) - Support Hive operation ALTER TABLE EXCHANGE.
- [SENTRY-2269](#) - Make SentryStore pluggable
- [SENTRY-2299](#) - NPE In Sentry HDFS Sync Plugin
- [SENTRY-2310](#) - Sentry is not be able to fetch full update subsequently, when there is HMS restart in the snapshot process

Apache Solr

- [SOLR-12290](#) - Do not close any servlet streams and improve our servlet stream closing prevention code for users and devs.
- [SOLR-12293](#) - Updates need to use their own connection pool to maintain connection reuse and prevent spurious recoveries.

Apache Spark

- [SPARK-12504](#) - [SQL] Masking credentials in the sql plan explain output for JDBC data sources
- [SPARK-12652](#) - [PYSPARK] Upgrade Py4J to 0.9.1
- [SPARK-13709](#) - [SQL] Initialize deserializer with both table and partition properties when reading partitioned tables
- [SPARK-13807](#) - De-duplicate `Python*Helper` instantiation code in PySpark streaming
- [SPARK-13848](#) - [SPARK-5185] Update to Py4J 0.9.2 in order to fix classloading issue
- [SPARK-15061](#) - [PYSPARK] Upgrade to Py4J 0.10.1
- [SPARK-16781](#) - [PYSPARK] java launched by PySpark as gateway may not be the same java used in the spark environment
- [SPARK-17960](#) - [PYSPARK][UPGRADE TO PY4J 0.10.4]
- [SPARK-19822](#) - [TEST] CheckpointSuite.testCheckpointedOperation: should not filter checkpointFilesOfLatestTime with the PATH string
- [SPARK-20862](#) - [MLLIB][PYTHON] Avoid passing float to ndarray.reshape in LogisticRegressionModel
- [SPARK-21278](#) - [PYSPARK] Upgrade to Py4J 0.10.6
- [SPARK-22429](#) - [STREAMING] Streaming checkpointing code does not retry after failure
- [SPARK-23852](#) - [SQL] Add test that fails if PARQUET-1217 is not fixed

Apache Sqoop

- [SQOOP-2567](#) - SQOOP import for Oracle fails with invalid precision/scale for decimal
- [SQOOP-3082](#) - Sqoop import fails after TCP connection reset if split by datetime column

Apache Zookeeper

- [ZOOKEEPER-2375](#) - Prevent multiple initialization of login object in each ZooKeeperSaslClient instance

Issues Fixed in CDH 5.15.0

CDH 5.15.0 fixes the following issues.

Impala/Sentry security roles mismatch after Catalog Server restart

This issue occurs when Impala's Catalog Server is restarted without also restarting all the Impala Daemons.

Impala uses generated numeric identifiers for roles. These identifiers are regenerated during catalogd restarts, and the same role can get a different identifier, possibly used by a different role before restart. An impalad's metadata cache can contain old id + role pairs, and when it is updated with privileges with new role ids from the catalog, the privilege will be added to the wrong role; the one that previously had the same role id.

Products affected: Apache Impala

Releases affected:

- CDH 5.14.4 and all prior releases

Users affected: Impala users with authorization enabled.

Date/time of detection: 5th October, 2018

Severity (Low/Medium/High): 3.8 "Low"; [CVSS](#):3.0/AV:N/AC:H/PR:H/UI:R/S:U/C:L/I:L/A:L

Impact: Users may get privileges of unrelated users.

CVE: CVE-2019-16381

Immediate action required: Update to a version of CDH containing the fix.

Addressed in release/refresh/patch: CDH 5.15.0

Apache Hive vulnerabilities CVE-2018-1282 and CVE-2018-1284

This security bulletin, [TSB 2018-299](#), covers two vulnerabilities that have been addressed in CDH 5.15.0 and later:

- **CVE-2018-1282:** JDBC driver is susceptible to SQL injection attack if the input parameters are not properly cleaned

- **CVE-2018-1284:** Hive UDF series `UDFXPathXXXX` allows users to pass carefully crafted XML to access arbitrary files

For further details, see the original [Known Issue for Hive](#).

Apache Oozie Server vulnerability CVE-2017-15712

A vulnerability in the Oozie Server allows a cluster user to read private files owned by the user running the Oozie Server process.

Products affected: Oozie

Releases affected: All releases prior to CDH 5.12.0, CDH 5.12.0, CDH 5.12.1, CDH 5.12.2, CDH 5.13.0, CDH 5.13.1, CDH 5.14.0

Users affected: Users running the Oozie Server

Date/time of detection: November 13, 2017

Detected by: Daryn Sharp and Jason Lowe of Oath (formerly Yahoo! Inc)

Severity: (Low/Medium/High) High

Impact: The vulnerability allows a cluster user to read private files owned by the user running the Oozie Server process. The malicious user can construct a workflow XML file containing XML directives and configuration that reference sensitive files on the Oozie server host.

CVE: CVE-2017-15712

Immediate action required: Upgrade to release where the issue is fixed.

Addressed in release/refresh/patch: CDH 5.13.2 and higher, 5.14.2 and higher, 5.15.0 and higher

Apache Hadoop MapReduce Job History Server (JHS) vulnerability CVE-2017-15713

A vulnerability in Hadoop's Job History Server allows a cluster user to expose private files owned by the user running the MapReduce Job History Server (JHS) process. See <http://seclists.org/oss-sec/2018/q1/79> for reference.

Products affected: Apache Hadoop MapReduce

Releases affected: All releases prior to CDH 5.12.0. CDH 5.12.0, CDH 5.12.1, CDH 5.12.2, CDH 5.13.0, CDH 5.13.1, CDH 5.14.0

Users affected: Users running the MapReduce Job History Server (JHS) daemon

Date/time of detection: November 8, 2017

Detected by: Man Yue Mo of lgtm.com

Severity (Low/Medium/High): High

Impact: The vulnerability allows a cluster user to expose private files owned by the user running the MapReduce Job History Server (JHS) process. The malicious user can construct a configuration file containing XML directives that reference sensitive files on the MapReduce Job History Server (JHS) host.

CVE: CVE-2017-15713

Immediate action required: Upgrade to a release where the issue is fixed.

Addressed in release/refresh/patch: CDH 5.13.2 and higher, 5.14.2 and higher, 5.15.0 and higher

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.15.0:

Apache Avro

- [AVRO-2109](#) - Reset buffers in case of IOException

Apache Flume

- [FLUME-2442](#) - Need an alternative to providing clear text passwords in flume config

- [FLUME-2917](#) - Provide netcat UDP source as alternative to TCP

Apache Hadoop

- [HADOOP-12886](#) - Exclude weak ciphers in SSLFactory through ssl-server.xml
- [HADOOP-14966](#) - Handle JDK-8071638 for hadoop-common
- [HADOOP-15085](#) - Output streams closed with IOUtils suppressing write errors
- [HADOOP-15113](#) - NPE in S3A getFileStatus: null instrumentation on using closed instance.
- [HADOOP-15149](#) - CryptoOutputStream should implement StreamCapabilities.
- [HADOOP-15161](#) - Streaming and common statistics missing from S3A metrics
- [HADOOP-15185](#) - Update ADLS connector to use the current version of ADLS SDK.
- [HADOOP-15206](#) - BZip2 drops and duplicates records when input split size is small
- [HDFS-1172](#) - Blocks in newly completed files are considered under-replicated too quickly
- [HDFS-7764](#) - DirectoryScanner shouldn't abort the scan if one directory had an error
- [HDFS-8693](#) - refreshNamenodes does not support adding a new standby to a running DN
- [HDFS-9023](#) - When NN is not able to identify DN for replication, reason behind it can be logged.
- [HDFS-10453](#) - ReplicationMonitor thread could stuck for long time due to the race between replication and delete of same file in a large cluster.
- [HDFS-10690](#) - Optimize insertion/removal of replica in ShortCircuitCache
- [HDFS-11187](#) - Optimize disk access for last partial chunk checksum of Finalized replica
- [HDFS-11494](#) - Log message when DN is not selected for block replication
- [HDFS-11576](#) - Block recovery will fail indefinitely if recovery time > heartbeat interval
- [HDFS-11847](#) - Enhance dfsadmin listOpenFiles command to list files blocking datanode decommissioning.
- [HDFS-11848](#) - Enhance dfsadmin listOpenFiles command to list files under a given path
- [HDFS-12318](#) - Fix IOException condition for openInfo in DFSInputStream
- [HDFS-12323](#) - NameNode terminates after full GC thinking QJM unresponsive if full GC is much longer than timeout
- [HDFS-12832](#) - INode.getFullPathName may throw ArrayIndexOutOfBoundsException lead to NameNode exit.
- [HDFS-12881](#) - Output streams closed with IOUtils suppressing write errors
- [HDFS-12910](#) - Secure Datanode Starter should log the port when it fails to bind
- [HDFS-13112](#) - Token expiration edits may cause log corruption or deadlock
- [HDFS-13170](#) - Port webhdfs unmaskedpermission parameter to HTTPFS
- [MAPREDUCE-6711](#) - JobImpl fails to handle preemption events on state COMMITTING

Apache HBase

Code Changes Might Be Required

The following fixed issues might require changes to your HBase code or your configuration:

- [HBASE-15128](#) - Disable region splits and merges switch in master.
This might require code changes if `org.apache.hadoop.hbase.client.Admin` was subclassed.
- [HBASE-15866](#) - Split `hbase.rpc.timeout` into `*.read.timeout` and `*.write.timeout`
This might require code changes if you want to specify different values for RPC read and write timeouts.
- [HBASE-16008](#) - A robust way deal with early termination of HBCK.
This might require code changes if `org.apache.hadoop.hbase.client.Admin` was subclassed.

Code Changes Should Not Be Required

The following fixes should not require code changes, but they contain improvements that might enhance your deployment:

- [HBASE-14252](#) - RegionServers fail to start when setting `hbase.ipc.server.callqueue.scan.ratio` to 0
- [HBASE-19163](#) - "Maximum lock count exceeded" from region server's batch processing
- [HBASE-19440](#) - Not able to enable balancer with RSGroups once disabled

- [HBASE-19886](#) - Display maintenance mode in shell, web UI

Apache Hive

Code Changes Might Be Required

The following fixed issues might require changes to your HiveQL code or your configuration:

- [HIVE-16324](#) - Truncate table should not work when EXTERNAL property of table is set to `true` (lower case). This fix changes how the EXTERNAL table property is interpreted by Hive. After this fix, if your current deployment sets the EXTERNAL table property to `true` (lower case), the external table will no longer be truncated. Read the Jira for further details.
- [HIVE-18879](#) - Disallow embedded elements in the UDFXPathUtil class. This fix might change how your XML parser works if you use embedded elements with the `xpath` UDFs. For example, if you use embedded elements with the following UDFs in your queries: `xpath`, `xpath_short`, `xpath_int`, `xpath_long`, `xpath_float`, `xpath_double`, `xpath_number`, or `xpath_string`.

Code Changes Should Not Be Required

The following fixes should not require code changes, but they contain improvements that might enhance your deployment:

- [HIVE-8472](#) - Add SET LOCATION option to the ALTER DATABASE command
- [HIVE-10495](#) - Hive index creation code throws a NullPointerException if index table is null
- [HIVE-14786](#) - Option added to Beeline to display binary column data as a string or a byte array (`--convertBinaryArrayToString=[true | false]`)
- [HIVE-14792](#) - Optimization added to minimize AvroSerde reads of the remote schema-file
- [HIVE-15329](#) - Fix for when a NullPointerException occurs during table creation
- [HIVE-15543](#) - Fix to prevent Hive fetching Spark memory/cores to decide parallelism when Spark dynamic allocation is enabled
- [HIVE-16601](#) - Display Session Id and Query Name / Id in Spark UI
- [HIVE-16663](#) - Added string caching for rows to Beeline
- [HIVE-16890](#) - Remove the superfluous wrapper from HiveVarcharWritable
- [HIVE-18228](#) - Azure credential properties added to the HiveConf hidden list
- [HIVE-18788](#) - Cleans up inputs in JDBC PreparedStatement (`HivePreparedStatement`) to fix SQL injection vulnerabilities

Hue

- [HUE-7860](#) - [core] Integrate non IO blocking Python Webserver
- [HUE-7913](#) - [autocomplete] Add variable locations to the autocomplete parser
- [HUE-7913](#) - [autocomplete] Report quoted variable locations with possible column references
- [HUE-7915](#) - [assist] Increase the frontend cache TTL to 10 days
- [HUE-7942](#) - [editor] FIX variables with incorrect placeholder.
- [HUE-7943](#) - [editor] Variable list are not being refreshed
- [HUE-8043](#) - [editor] Cancel a query using ctrl-enter shortcut

Apache Impala

- [IMPALA-4315](#) - Allow `USE` and `SHOW TABLES` if there is at least one table in a database where the user has table or column privileges.
- [IMPALA-4323](#) - The `SET ROW FORMAT` clause was added to the `ALTER TABLE` statement for the `TEXT` or `SEQUENCE` file formats.
- [IMPALA-4886](#) - Table metrics are available in the catalog web UI.
- [IMPALA-5654](#) - Disallows explicitly setting the Kudu table name property for managed Kudu tables in a `CREATE TABLE` and `ALTER TABLE` statements, e.g. `CREATE TABLE t (i INT) STORED AS KUDU TBLPROPERTIES('kudu.table_name' = 'some_name')`.
- [IMPALA-6549](#) - The file handle cache, controlled by the `max_cached_file_handles` flag, is enabled by default.

Apache Kudu

- [KUDU-1613](#) - Fixed a scenario where the on-disk data of a tablet server was completely erased and a new tablet server was started on the same host. This issue could prevent tablet replicas previously hosted on the server from being evicted and re-replicated. Tablets now immediately evict replicas that respond with a different server UUID than expected.
- [KUDU-1927](#) - Fixed a rare race condition when connecting to masters during their startup which might cause a client to get a response without a CA certificate and/or authentication token. This would cause the client to fail to authenticate with other servers in the cluster. The leader master now always sends a CA certificate and an authentication token (when applicable) to a Kudu client with a successful ConnectToMaster response.
- [KUDU-2262](#) - The Kudu Java client now will retry a connection if no master is discovered as a leader, and the user has a valid authentication token. This avoids failure in recoverable cases when masters are in the process of the very first leader election after starting up.
- [KUDU-2264](#) - The Java client will now automatically attempt to re-acquire Kerberos credentials from the ticket cache when the prior credentials are about to expire. This allows client instances to persist longer than the expiration time of a single Kerberos ticket so long as some other process renews the credentials in the ticket cache. Documentation on interacting with Kerberos authentication has been added to the Javadoc for the `AsyncKuduClient` class.
- [KUDU-2265](#) - Follower masters are now able to verify authentication tokens even if they have never been a leader. Prior to this fix, if a follower master had never been a leader, clients would be unable to authenticate to that master, resulting in spurious error messages being logged.
- [KUDU-2295](#) - Fixed a tablet server crash when a tablet replica is deleted during a scan.
- [KUDU-2312](#) - The evaluation order of predicates in scans with multiple predicates has been made deterministic. Due to a bug, this was not necessarily the case previously. Predicates are applied in most to least selective order, with ties broken by column index. The evaluation order may change in the future, particularly when better column statistics are made available internally.
- [KUDU-2331](#) - Previously, the `kudu tablet change_config move_replica` tool required all tablet servers in the cluster to be available when performing a move. This restriction has been relaxed: only the tablet server that will receive a replica of the tablet being moved and the hosts of the tablet's existing replicas need to be available for the move to occur.
- [KUDU-2343](#) - Fixed a bug in the Java client which prevented the client from locating the new leader master after a leader failover in the case that the previous leader either remained online or restarted quickly. This bug resulted in the client timing out operations with errors indicating that there was no leader master.
- [KUDU-2259](#) - The Unix process username of the client is now included inside the exported security credentials, so that the effective username of clients who import credentials and subsequently use unauthenticated (SASL PLAIN) connections matches the client who exported the security credentials. For example, this is useful to let the Spark executors know which username to use if the Spark driver has no authentication token. This change only affects clusters with encryption disabled using `--rpc-encryption=disabled`.

Apache Oozie

- [OOZIE-3173](#) - Coordinator job with frequency using cron syntax creates only one action in catchup mode
- [OOZIE-3183](#) - Better logging for SshActionExecutor and extended HA capability when calling to remote host

Apache Spark

- [SPARK-12297](#) - Convert Impala-written timestamps from UTC to local TZ ()
- [SPARK-22188](#) - [CORE] Adding security headers for preventing XSS, MitM and MIME sniffing
- [SPARK-23660](#) - Fix exception in yarn cluster mode when application ended fast

Apache Sqoop

- [SQOOP-3153](#) - Sqoop export with `--as-<spec_file_format>` will now display a verbose error message as these options are only valid for imports

Issues Fixed in CDH 5.14.x

The following topics describe issues fixed in CDH 5.14.x, from newest to oldest release. You can also review [What's New in CDH 5.14.x](#) on page 1816 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.14.4

CDH 5.14.4 fixes the following issues.

Hadoop YARN Privilege Escalation CVE-2016-6811

Fixed a vulnerability in Hadoop YARN that allows a user who can escalate to the yarn user the ability to possibly run arbitrary commands as the root user. CVE-2016-6811

Apache HBase potential privilege escalation for user of HBase “Thrift 1” API Server over HTTP CVE-2018-8025

CVE-2018-8025 describes an issue in Apache HBase that affects the optional "Thrift 1" API server when running over HTTP. There is a race-condition that could lead to authenticated sessions being incorrectly applied to users, e.g. one authenticated user would be considered a different user or an unauthenticated user would be treated as an authenticated user.

Products affected: HBase Thrift Server

Releases affected:

- CDH 5.4.x - 5.12.x
- CDH 5.13.0, 5.13.1, 5.13.2, 5.13.3
- CDH 5.14.0, 5.14.2, 5.14.3
- CDH 5.15.0

Fixed versions: CDH 5.14.4

Users affected: Users with the HBase Thrift 1 service role installed and configured to work in “thrift over HTTP” mode. For example, those using Hue with HBase impersonation enabled.

Severity: High

Potential privilege escalation.

CVE: CVE-2018-8025

Immediate action required: Upgrade to a CDH version with the fix, or, disable the HBase Thrift-over-HTTP service. Disabling the HBase Thrift-over-HTTP service will render Hue impersonation inoperable and all HBase access via Hue will be performed using the “hue” user instead of the authenticated user.

Knowledge article: For the latest update on this issue see the corresponding Knowledge article - TSB: 2018-315: Potential privilege escalation for user of HBase “Thrift 1” API Server over HTTP

Zip Slip Vulnerability CVE-2018-8009

“Zip Slip” is a widespread arbitrary file overwrite critical vulnerability, which typically results in remote command execution. It was discovered and responsibly disclosed by the Snyk Security team ahead of a public disclosure on June 5, 2018, and affects thousands of projects.

Cloudera has analyzed our use of zip-related software, and has determined that only Apache Hadoop is vulnerable to this class of vulnerability in CDH 5. This has been fixed in upstream Hadoop as CVE-2018-8009.

Apache HiveQL SELECT STRUCT fails with RuntimeException

The HiveQL query `SELECT STRUCT("x" , null) ;` fails in CDH 5.10.0 onward and returns the following compilation error:

```
Error while compiling statement: FAILED: RuntimeException
java.lang.RuntimeException: Hive internal error: conversion of void to void not
supported yet.
```


This has been fixed in CDH 5.14.4 (CDH-67208) and later and now HiveQL queries like `SELECT STRUCT("x" ,null) ;` run successfully.

Apache Hive vulnerabilities [CVE-2018-1282](#) and [CVE-2018-1284](#)

This security bulletin, [TSB 2018-299](#), covers two vulnerabilities that have been addressed in CDH 5.14.4 and later:

- **CVE-2018-1282:** JDBC driver is susceptible to SQL injection attack if the input parameters are not properly cleaned
- **CVE-2018-1284:** Hive UDF series `UDFXPathXXXX` allows users to pass carefully crafted XML to access arbitrary files

For further details, see the original [Known Issue for Hive](#).

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.14.4:

Apache Flume

- [FLUME-2786](#) - FLUME-3056 FLUME-3117 Application enters a deadlock when stopped while `handleConfigurationEvent`
- [FLUME-2894](#) - Flume components should stop in the correct order

Apache Hadoop

- [HADOOP-12886](#) - Exclude weak ciphers in SSLFactory through `ssl-server.xml`
- [HADOOP-13024](#) - Distcp with `-delete` feature on raw data not implemented
- [HADOOP-13972](#) - ADLS to support per-store configuration.
- [HADOOP-15186](#) - Allow Azure Data Lake SDK dependency version to be set on the command line.
- [HADOOP-15317](#) - Improve `NetworkTopology.chooseRandom`'s loop.
- [HADOOP-15342](#) - Update ADLS connector to use the current SDK version (2.2.7)
- [HADOOP-15356](#) - Make HTTP timeout configurable in ADLS connector
- [HADOOP-15466](#) - Correct units in `adl.http.timeout`.
- [HDFS-9229](#) - Expose size of NameNode directory as a metric
- [HDFS-10903](#) - Replace config key literal strings with config key names II: `hadoop hdfs`
- [HDFS-11993](#) - Add log info when connect to datanode socket address failed
- [HDFS-12832](#) - `INode.getFullPathName` may throw `ArrayIndexOutOfBoundsException` lead to NameNode exit.
- [HDFS-13170](#) - Port `webhdfs.unmaskedpermission` parameter to HTTPFS
- [HDFS-13314](#) - NameNode should optionally exit if it detects `FsImage` corruption
- [MAPREDUCE-7094](#) - `LocalDistributedCacheManager` leaves classloaders open, which leaks FDs
- [YARN-5121](#) - Fix some container-executor portability issues

Apache HBase

- [HBASE-15866](#) - Split `hbase.rpc.timeout` into `*.read.timeout` and `*.write.timeout`
- [HBASE-16201](#) - fix a NPE issue in `RpcServer`
- [HBASE-17944](#) - Removed unused JDK version parsing from `ClassSize`.
- [HBASE-20229](#) - `ConnectionImplementation.locateRegions()` returns duplicated entries when region replication is on
- [HBASE-20293](#) - `get_splits` returns duplicate split points when region replication is on
- [HBASE-20664](#) - Reduce the broad scope of `outToken` in `ThriftHttpServlet`

Apache Hive

Code Changes Might Be Required

No upstream fixed issues might require code changes for the CDH 5.14.4 release of Apache Hive.

Code Changes Should Not Be Required

The following fixes should not require code changes, but they contain improvements that might enhance your deployment:

- [HIVE-14792](#) - AvroSerde reads the remote schema-file at least once per mapper, per table reference.
- [HIVE-18788](#) - Clean up inputs in JDBC PreparedStatement
- [HIVE-18879](#) - Disallow embedded element in UDFXPathUtil needs to work if xercesImpl.jar in classpath
- [HIVE-19041](#) - Thrift deserialization of Partition objects should intern fields
- [HIVE-19424](#) - NPE In MetaDataFormatters
- [HIVE-19700](#) - Workaround for JLine issue with UnsupportedTerminal

Hue

- [HUE-7826](#) - [security] Uncaught TypeError: privilege.privilegeScope is not a function error when deleting a Solr privilege
- [HUE-7946](#) - [jib] Link to subworkflow on workflow dashboard page 404s in Hue 4
- [HUE-7946](#) - [jib] Link to subworkflow on workflow dashboard page 404s in Hue 4
- [HUE-7989](#) - [useradmin] Provide better UI message and message in logs when ldap server down
- [HUE-8063](#) - [fb] Better error message when HTTPFS is not working
- [HUE-8069](#) - [core] Pyformance reporter thread stops on exceptions in gauge callbacks
- [HUE-8159](#) - [oozie] Unable to create Workflow/Schedule using Java document Action
- [HUE-8170](#) - [useradmin] Fix LDAP sync (ldap_access.py) certificate validation logic
- [HUE-8173](#) - [core] Add a warning log when a users switched to the old Hue 3 UI
- [HUE-8173](#) - [core] Add a log line when a user loads a page not via the load balancer even if we have one
- [HUE-8191](#) - [jib] UnicodeEncodeError occurs in Job Browser when browser language is not English
- [HUE-8202](#) - [jib] Fix mutual authentication failed with Isilon
- [HUE-8207](#) - [indexer] Issue previewing input file with unicode data
- [HUE-8217](#) - [dashboard] Fix HTML resultset widget templating
- [HUE-8232](#) - [oozie] Fix 500 error if coordinator associated workflow has been deleted
- [HUE-8236](#) - [core] Correct Hue config value use integer instead of math expression
- [HUE-8274](#) - [s3] Moving a folder using drag and drop deletes the folder itself
- [HUE-8280](#) - [fb] Move action button does not prevent move to itself
- [HUE-8282](#) - [fb] Check if EC2 instance before check IAM metadata
- [HUE-8314](#) - [core] Fix SAML encryption missing config
- [HUE-8320](#) - [core] After copying a workflow using Hue 4 button, saving the copied workflow fails

Apache Impala

- [IMPALA-4315](#) - Users with the column-level privileges can run the USE and SHOW TABLES statements.
- [IMPALA-5270](#) - Impala now passes fully resolved sort expressions to prevent crashes with inline view and analytic functions.
- [IMPALA-6500](#) - Fixed an issue where Impala crashed under certain hypervisors that returned out-of-range CPU IDs.
- [IMPALA-6687](#) - The INSERT statements with mixed case partition column name no longer fails.
- [IMPALA-6822](#) - A new query option, SHUFFLE_DISTINCT_EXPRS, is available for disabling shuffling by distinct expressions.
- [IMPALA-6908](#) - The ECONNRESET error code handling was added to the IsConnReset function that checks if a given exception is due to a stale connection.
- [IMPALA-6934](#) - Correct results returns from the EXISTS subquery containing ORDER BY, LIMIT, and OFFSET
- [IMPALA-6954](#) - Fixed a problem with expression rewrites that caused losses of partitions with CTAS into Kudu.
- [IMPALA-7032](#) - Codegen for CHAR type null literals are disabled to prevent a crash.

Apache Kudu

- [KUDU-2331](#) - The `kudu tablet change_config move_replica` tool no longer fails for an unavailable tablet server if the tablet server is not the source or destination of the move.
- Improved the speed of scans in cases where a large number of rows have previously been deleted from the table.

Apache Oozie

- [OOZIE-2491](#) - oozie acl cannot specify group, it does not work

Apache Sentry

- [SENTRY-1209](#) - Sentry does not block Hive's cross-schema table renames
- [SENTRY-2144](#) - Table Rename Cross Database should update permission correctly.
- [SENTRY-2165](#) - NotificationProcessor process notification methods have logs wrongly flagged as ERROR
- [SENTRY-2183](#) - Increase default sentry-hdfs rpc timeout to 20 mins
- [SENTRY-2184](#) - Performance Issue: MPath is queried for each MAuthzPathsMapping in full snapshot
- [SENTRY-2194](#) - Upgrade Sentry hadoop-version dependency to 2.7.5
- [SENTRY-2226](#) - CDH5Support Hive operation ALTER TABLE EXCHANGE.

Apache Solr

- [SOLR-12293](#) - Updates need to use their own connection pool to maintain connection reuse and prevent spurious recoveries.

Apache Spark

- [SPARK-12504](#) - [SQL] Masking credentials in the sql plan explain output for JDBC data sources.
- [SPARK-12652](#) - [PYSPARK] Upgrade Py4J to 0.9.1
- [SPARK-13709](#) - [SQL] Initialize deserializer with both table and partition properties when reading partitioned tables
- [SPARK-13807](#) - De-duplicate `Python*Helper` instantiation code in PySpark streaming
- [SPARK-13848](#) - [SPARK-5185] Update to Py4J 0.9.2 in order to fix classloading issue
- [SPARK-15061](#) - [PYSPARK] Upgrade to Py4J 0.10.1
- [SPARK-16781](#) - [PYSPARK] java launched by PySpark as gateway may not be the same java used in the spark environment
- [SPARK-17960](#) - [PYSPARK][UPGRADE TO PY4J 0.10.4]
- [SPARK-20862](#) - [MLLIB][PYTHON] Avoid passing float to ndarray.reshape in LogisticRegressionModel
- [SPARK-21278](#) - [PYSPARK] Upgrade to Py4J 0.10.6
- [SPARK-23660](#) - Fix exception in yarn cluster mode when application ended fast

Apache Sqoop

- [SQOOP-2567](#) - SQOOP import for Oracle fails with invalid precision/scale for decimal
- [SQOOP-3082](#) - Sqoop import fails after TCP connection reset if split by datetime column

Issues Fixed in CDH 5.14.3

There are no fixed issues in CDH 5.14.3. This release only contained fixed issues for Cloudera Manager. For more information, see [Issues Fixed in Cloudera Manager 5.14.3](#) on page 2561.

Issues Fixed in CDH 5.14.2

CDH 5.14.2 fixes the following issues.

Missing results in Hive, Spark, Pig, Custom MapReduce jobs, and other Java applications when filtering Parquet data written by Impala

This issue, which was also reported in [TSB-2018-300](#), is fixed in CDH 5.14.2 and higher. For further details, see the original [Known Issue for Hive](#), [Known Issue for Spark](#), [Known Issue for Pig](#), or the [Known Issue for MapReduce](#).

Apache Oozie Server Vulnerability CVE-2017-15712

A vulnerability in the Oozie Server allows a cluster user to read private files owned by the user running the Oozie Server process.

Products affected: Oozie

Releases affected: All releases prior to CDH 5.12.0, CDH 5.12.0, CDH 5.12.1, CDH 5.12.2, CDH 5.13.0, CDH 5.13.1, CDH 5.14.0

Users affected: Users running the Oozie Server

Date/time of detection: November 13, 2017

Detected by: Daryn Sharp and Jason Lowe of Oath (formerly Yahoo! Inc)

Severity: (Low/Medium/High) High

Impact: The vulnerability allows a cluster user to read private files owned by the user running the Oozie Server process. The malicious user can construct a workflow XML file containing XML directives and configuration that reference sensitive files on the Oozie server host.

CVE: CVE-2017-15712

Immediate action required: Upgrade to release where the issue is fixed.

Addressed in release/refresh/patch: CDH 5.13.2, 5.14.2

[Apache Hadoop MapReduce Job History Server \(JHS\) vulnerability CVE-2017-15713](#)

A vulnerability in Hadoop's Job History Server allows a cluster user to expose private files owned by the user running the MapReduce Job History Server (JHS) process. See <http://seclists.org/oss-sec/2018/q1/79> for reference.

Products affected: Apache Hadoop MapReduce

Releases affected: All releases prior to CDH 5.12.0. CDH 5.12.0, CDH 5.12.1, CDH 5.12.2, CDH 5.13.0, CDH 5.13.1, CDH 5.14.0

Users affected: Users running the MapReduce Job History Server (JHS) daemon

Date/time of detection: November 8, 2017

Detected by: Man Yue Mo of lgtm.com

Severity (Low/Medium/High): High

Impact: The vulnerability allows a cluster user to expose private files owned by the user running the MapReduce Job History Server (JHS) process. The malicious user can construct a configuration file containing XML directives that reference sensitive files on the MapReduce Job History Server (JHS) host.

CVE: CVE-2017-15713

Immediate action required: Upgrade to a release where the issue is fixed.

Addressed in release/refresh/patch: CDH 5.13.2, 5.14.2, 5.15.0

[Slow or stalled queries under highly concurrent write workloads when Sentry + Hive are used, caused by a Hive metastore deadlock](#)

This issue, which was also reported in [TSB-2018-294](#), is fixed in CDH 5.14.2. For further details, see the original [Known Issue](#).

[Crash due to CHECK failure during compactions: pv_delete_redo != nullptr](#)

This issue, which was also reported in [TSB-2018-298](#), is fixed in CDH 5.13.3 and 5.14.2. For further details, see the original [Known Issue](#).

[Secure Impala-Kudu clusters need to be restarted frequently](#)

Due to Kudu-2264, secure, long-running Impala clusters on CDH 5.13.x and later will be unable to interact with Kudu tables after some time because the Kudu client used by Impala will no longer be able to authenticate with Kudu. The problem starts after the client's ticket can no longer be renewed, which is typically after 7 days. Restarting the Impala service will allow Impala to work with Kudu again.

Impala will begin failing to read or write Kudu tables with errors in the Impala daemon logs resembling:

```
W1121 22:47:47.231425 50416 ConnectToCluster.java:302] Error receiving response from
cdh-master-4db933ef.cdh-cluster.internal:7051
Java exception follows:
org.apache.kudu.client.NonRecoverableException: Server requires Kerberos, but this client
```

```

is not
authenticated (kinit)

at org.apache.kudu.client.Negotiator.evaluateChallenge(Negotiator.java:705)
at org.apache.kudu.client.Negotiator.sendSaslInitiate(Negotiator.java:581)
at org.apache.kudu.client.Negotiator.startAuthentication(Negotiator.java:545)
at org.apache.kudu.client.Negotiator.handleTlsMessage(Negotiator.java:499)
at org.apache.kudu.client.Negotiator.handleResponse(Negotiator.java:264)
at org.apache.kudu.client.Negotiator.messageReceived(Negotiator.java:231)
at ...
Caused by: javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException:
  No valid credentials provided (Mechanism level: Failed to find any Kerberos tgt)]

at com.sun.security.sasl.gsskerb.GssKrb5Client.evaluateChallenge(GssKrb5Client.java:211)
at org.apache.kudu.client.Negotiator$1.run(Negotiator.java:691)
at org.apache.kudu.client.Negotiator$1.run(Negotiator.java:688)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:422)
at org.apache.kudu.client.Negotiator.evaluateChallenge(Negotiator.java:687)
... 35 more

Caused by: GSSException: No valid credentials provided (Mechanism level: Failed to find
any
Kerberos tgt)

at sun.security.jgss.krb5.Krb5InitCredential.getInstance(Krb5InitCredential.java:147)
at sun.security.jgss.krb5.Krb5MechFactory.getCredentialElement(Krb5MechFactory.java:122)
at sun.security.jgss.krb5.Krb5MechFactory.getMechanismContext(Krb5MechFactory.java:187)
at sun.security.jgss.GSSManagerImpl.getMechanismContext(GSSManagerImpl.java:224)
at sun.security.jgss.GSSContextImpl.initSecContext(GSSContextImpl.java:212)
at sun.security.jgss.GSSContextImpl.initSecContext(GSSContextImpl.java:179)
at com.sun.security.sasl.gsskerb.GssKrb5Client.evaluateChallenge(GssKrb5Client.java:192)
... 40 more

```

Alternatively, errors may be seen like the following:

```

I0102 11:58:36.688992 3587 jni-util.cc:196]
org.apache.impala.common.ImpalaRuntimeException: Unable to initialize the Kudu scan node

at org.apache.impala.planner.KuduScanNode.init(KuduScanNode.java:133)
at org.apache.impala.planner.SingleNodePlanner.createScanNode(SingleNodePlanner.java:1312)
at
org.apache.impala.planner.SingleNodePlanner.createTableRefNode(SingleNodePlanner.java:1519)
at
org.apache.impala.planner.SingleNodePlanner.createTableRefsPlan(SingleNodePlanner.java:779)
at
org.apache.impala.planner.SingleNodePlanner.createSelectPlan(SingleNodePlanner.java:617)
at org.apache.impala.planner.SingleNodePlanner.createQueryPlan(SingleNodePlanner.java:260)
at
org.apache.impala.planner.SingleNodePlanner.createSingleNodePlan(SingleNodePlanner.java:150)
at org.apache.impala.planner.Planner.createPlan(Planner.java:98)
at org.apache.impala.service.Frontend.createExecRequest(Frontend.java:1004)
at org.apache.impala.service.Frontend.createExecRequest(Frontend.java:1100)
at org.apache.impala.service.JniFrontend.createExecRequest(JniFrontend.java:156)

Caused by: org.apache.kudu.client.NonRecoverableException: cannot re-acquire
authentication
token after 5 attempts

at
org.apache.kudu.client.AuthnTokenReacquirer$1NewAuthnTokenErrB.failQueuedRpcs(AuthnTokenReacquirer.java:167)
at
org.apache.kudu.client.AuthnTokenReacquirer$1NewAuthnTokenErrB.call(AuthnTokenReacquirer.java:159)
at
org.apache.kudu.client.AuthnTokenReacquirer$1NewAuthnTokenErrB.call(AuthnTokenReacquirer.java:142)
at com.stumbleupon.async.Deferred.doCall(Deferred.java:1262)
at com.stumbleupon.async.Deferred.runCallbacks(Deferred.java:1241)
at com.stumbleupon.async.Deferred.callback(Deferred.java:989)
at
org.apache.kudu.client.ConnectToCluster.incrementCountAndCheckExhausted(ConnectToCluster.java:223)
at org.apache.kudu.client.ConnectToCluster.access$000(ConnectToCluster.java:48)
at

```

```
org.apache.kudu.client.ConnectToCluster$ConnectToMasterErrCB.call(ConnectToCluster.java:304)
at
org.apache.kudu.client.ConnectToCluster$ConnectToMasterErrCB.call(ConnectToCluster.java:293)
at com.stumbleupon.async.Deferred.doCall(Deferred.java:1262)
at com.stumbleupon.async.Deferred.runCallbacks(Deferred.java:1241)
at com.stumbleupon.async.Deferred.handleContinuation(Deferred.java:1297)
```

Products Affected: Impala

Releases Affected: CDH 5.13.0, CDH 5.13.1, and CDH 5.14.0

Users Affected: Users running a secure cluster with Impala and Kudu.

Severity: Medium

Impact: Impala will not be able to query or write to Kudu tables until the service is restarted. The Impala service will need to be restarted periodically in order to continue being able to work with Kudu tables.

Immediate action required: To work around the issue, users can restart the Impala service at least once per ticket renewal period. This period is typically 7 days.

Addressed in release: CDH 5.13.2, CDH 5.13.3, and CDH 5.14.2

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.14.2:

- [HADOOP-6852](#) - apparent bug in concatenated-bzip2 support (decoding)
- [HADOOP-14872](#) - CryptoInputStream should implement unbuffer
- [HADOOP-15012](#) - Add readahead, dropbehind, and unbuffer to StreamCapabilities
- [HADOOP-15143](#) - NPE due to Invalid KerberosTicket in UGI
- [HDFS-4185](#) - Add a metric for number of active leases
- [HDFS-6054](#) - MiniQJMHACluster should not use static port to avoid binding failure in unit test.
- [HDFS-11644](#) - Support for querying outputstream capabilities
- [HDFS-12251](#) - Add document for StreamCapabilities.
- [HDFS-12386](#) - Add fsserver defaults call to WebhdfsFileSystem.
- [HDFS-12528](#) - Added a configuration dfs.domain.socket.disable.interval.seconds to change the interval that disables short-circuit reads on failures.
- [HDFS-12985](#) - NameNode crashes during restart after an OpenForWrite file present in the Snapshot got deleted.
- [HDFS-13040](#) - Kerberized inotify client fails despite kinit properly
- [HDFS-13115](#) - In getNumUnderConstructionBlocks(), ignore the inodelds for which the inodes have been deleted
- [HDFS-13164](#) - File not closed if streamer fail with DSQuotaExceededException.
- [MAPREDUCE-7015](#) - Possible race condition in JHS if the job is not loaded
- [MAPREDUCE-7020](#) - Task timeout in uber mode can crash AM
- [MAPREDUCE-7048](#) - Uber AM can crash due to unknown task in statusUpdate
- [MAPREDUCE-7053](#) - Timed out tasks can fail to produce thread dump
- [YARN-2416](#) - InvalidStateTransitionException in ResourceManager if AMLauncher does not receive response for startContainers() call in time
- [YARN-3760](#) - FSDataOutputStream leak in AggregatedLogFormat.LogWriter.close()
- [YARN-4813](#) - TestRMWebServicesDelegationTokenAuthentication.testDoAs fails intermittently
- [YARN-6288](#) - Exceptions during aggregated log writes are mishandled
- [YARN-7083](#) - Log aggregation deletes/renames while file is open
- [YARN-7363](#) - ContainerLocalizer don't have a valid log4j config in case of Linux container executor.
- [YARN-7590](#) - Improve container-executor validation check
- [YARN-7622](#) - Allow fair-scheduler configuration on HDFS
- [YARN-7705](#) - Create the container log directory with correct sticky bit in C code
- [YARN-7758](#) - Add an additional check to the validity of container and application ids passed to container-executor
- [HBASE-14252](#) - RegionServers fail to start when setting hbase.ipc.server.callqueue.scan.ratio to 0

- [HBASE-16980](#) - TestRowProcessorEndpoint failing consistently
- [HBASE-17924](#) - Consider sorting the row order when processing multi() ops before taking rowlocks
- [HBASE-18625](#) - Splitting of region with replica, doesn't update region list in serverHolding. A server crash leads to overlap.
- [HBASE-19163](#) - "Maximum lock count exceeded" from region server's batch processing
- [HBASE-19364](#) - Truncate_preserve fails with table when replica region > 1
- [HBASE-19391](#) - Calling HRegion#initializeRegionInternals from a region replica can still re-create a region directory
- [HBASE-19553](#) - Old replica regions should be cleared from AM memory after primary region split or merge
- [HBASE-19728](#) - Add lock to filesCompacting in all place.
- [HBASE-19934](#) - HBaseSnapshotException when read replicas is enabled and online snapshot is taken after region splitting
- [HBASE-20061](#) - HStore synchronized member variable filesCompacting should be private
- [HIVE-11569](#) - Use PreOrderOnceWalker where feasible
- [HIVE-11609](#) - Capability to add a filter to hbase scan via composite key doesn't work
- [HIVE-14792](#) - AvroSerde reads the remote schema-file at least once per mapper, per table reference.
- [HIVE-15859](#) - HoS: Write RPC messages in event loop
- [HIVE-17624](#) - MapredLocalTask running in separate JVM could throw ClassNotFoundException
- [HIVE-18127](#) - Do not strip '--' comments from shell commands issued from CliDriver
- [HIVE-16886](#) - Hive metastore (HMS) log notifications may have duplicated event IDs if multiple HMSs are running concurrently
- [HIVE-18671](#) - Lock not released after Hive on Spark query was cancelled
- [HIVE-18917](#) - Add spark.home to hive.conf.restricted.list
- [HUE-7646](#) - [sentry] Changing role privileges removes HDFS ACLs
- [HUE-7755](#) - [oozie] Adding Distcp arguments and properties
- [HUE-7779](#) - [jb] Add job type in the API calls
- [HUE-7783](#) - [editor] Add log toggling functionality to the presentation mode
- [HUE-7784](#) - [dashboard] Fix JS error on right assist filtering
- [HUE-7790](#) - [assist] Silence all API errors in the right assistant
- [HUE-7794](#) - [fb] Modify/upload file failed if S3 bucket encryption is SSE-KMS
- [HUE-7802](#) - [dashboard] Protect against missing function in old dashboards
- [HUE-7806](#) - [core] Show all the users "Switch to Hue 4" option
- [HUE-7812](#) - [core] Metrics and Threads page don't get refreshed Step#1
- [HUE-7813](#) - [useradmin] Add login dropdown when there is only one item
- [HUE-7822](#) - [core] Unable to clone oozie workflow in Hue 4.0
- [HUE-7830](#) - [core] Non supported feature use_default_configuration doesn't work in Hue 4
- [HUE-7837](#) - [presentation] Toggle hiding/showing the code snippets should also hide the log panels
- [HUE-7842](#) - [dashboard] Avoid extra search call when dropping a new facet in a new row
- [HUE-7844](#) - [editor] Fix rendering of barcharts when the first value is missing
- [HUE-7850](#) - [editor] Presentation mode will fail if past queries are expired
- [HUE-7854](#) - [core] Update translations files for 4.2
- [HUE-7855](#) - [editor] Fixed JS error on subscription removal
- [HUE-7855](#) - [editor] Disable the editor auto-resize when running a saved query
- [HUE-7861](#) - [jb] Handle Impala query timeout
- [HUE-7874](#) - [oozie] SSH action missing configuration for retry-max and retry-interval
- [HUE-7898](#) - [editor] Impala query expiration shows update max_row_size help
- [HUE-7909](#) - [editor] Hide options in session settings
- [HUE-7916](#) - [frontend] Avoid breaking of scatterplot chart for some grouping values
- [HUE-7920](#) - [frontend] The stacked graph shouldn't have empty areas in the stacks
- [HUE-7927](#) - [metastore] Sample view on the main table page should have an horizontal scrollbar
- [HUE-7929](#) - [indexes] Field list table on the index page should have an horizontal scrollbar

- [HUE-7932](#) - [oozie] Clicking on workflow with several decision nodes hangs Hue
- [HUE-7934](#) - [importer] Unselecting has Header checkbox has no effect
- [HUE-7974](#) - [metastore] Open Database location opens in Hue 3
- [HUE-7992](#) - [oozie] Big workflows show horizontally scroll in Hue 4
- [HUE-7997](#) - [backend] Python LDAP library exposes password
- [HUE-8009](#) - [useradmin] Fix login page error for invalid local user when LDAP also enabled
- [HUE-8024](#) - [indexer] Use old collection API to avoid grant Update privilege
- [HUE-8025](#) - [metadata] Do not show the deleted databasess/tables in the search
- [HUE-8032](#) - [core] XSS in Hue Admin Interface
- [HUE-8036](#) - [presentation] Protect from storm of API calls when loading back a query with presentation
- [HUE-8053](#) - [useradmin] LDAP authentication with sync_groups_on_login=true fails with KeyError exception
- [HUE-8055](#) - [desktop] Support multiple LDAP servers in LDAP Test command
- [IMPALA-3942](#) - Fixed incorrectly escaped string literal in front-end.
- [IMPALA-4664](#) - Fixed an unexpected string conversion in Shell.
- [IMPALA-5139](#) - Updated mvn-quiet.sh to print execution content to log file.
- [IMPALA-5269](#) - Fix the issue with final line of query followed by a comment.
- [IMPALA-5936](#) - The '%' operator no longer overflows on large decimals.
- [IMPALA-5940](#) - Impala no longer shows stack tracing with Status::Expected().
- [IMPALA-6008](#) - Creating a UDF from a shared library with a .ll extension does not crash Impala.
- [IMPALA-6114](#) - Require type equality for NumericLiteral::localEquals().
- [IMPALA-6284](#) - Mark the intermediate decimal avg struct as packed.
- [IMPALA-6295](#) - Fix mix/max handling of 'nan' and 'inf'.
- [IMPALA-6307](#) - CTAS statement fails with duplicate column exception.
- [IMPALA-6348](#) - Redact only sensitive fields in runtime profiles.
- [IMPALA-6353](#) - Impala no longer crashes in snappy decompressor.
- [IMPALA-6384](#) - RequestPoolService should honor custom group mapping config.
- [IMPALA-6388](#) - Fix the Union node number of hosts estimation
- [IMPALA-6418](#) - Find a reliable way to detect supported TLS versions.
- [IMPALA-6435](#) - Disabled codegen for CHAR literals.
- [IMPALA-6454](#) - CTAS into Kudu fails with mixed-case partition or primary key column names.
- [IMPALA-6473](#) - An analytic sorts with the same expression in PARTITION BY and ORDER BY does not return an error.
- [IMPALA-6489](#) - Impala now uses the correct template tuple size.
- [IMPALA-6530](#) - Impala now tracks time spent opening HDFS file handles.
- [IMPALA-6538](#) - Fixed read path when Parquet min/max statistics contain NaN.
- [IMPALA-6542](#) - Fixed inconsistent write path of Parquet min/max statistics.
- [IMPALA-6570](#) - Support KUDU on SLES12.
- [IMPALA-6577](#) - Avoid slow Status constructor on expiration thread.
- [KUDU-1613](#) - Fixed an issue where a reformatted server could prevent re-replication of the replicas it previously hosted.
- [KUDU-2238](#) - Fixed an issue where large updates were not flushed to disk under memory pressure.
- [KUDU-2251](#) - Fixed a crash due to an overflow bug that could occur when a tablet was updated very frequently.
- [KUDU-2274](#) - Fixed an issue where copying over a tombstoned replica could produce a replica in an inconsistent state.
- [KUDU-2343](#) - Fixed an issue where the Java client would fail to connect to the leader master after a leadership change.
- [OOZIE-1401](#) - amend PurgeCommand should purge the workflow jobs w/o end_time
- [OOZIE-2045](#) - Oozie can detect symlink changes and reload the sharelib
- [OOZIE-2536](#) - Hadoop's cleanup of local directory in uber mode causing failures

- [PARQUET-1133](#) - Add int96 support by returning bytearray, Skip originalType comparison for map types when originalType is null
- [PARQUET-1217](#) - Incorrect handling of missing values in Statistics
- [PIG-5325](#) - Schema disambiguation can't be turned off for nested schemas
- [SENTRY-1459](#) - Alter view with HMS Client fails with "java.lang.IllegalArgumentException: Can not create a Path from a null string"
- [SENTRY-1951](#) - Old SentryStore.retrieveFullPathsImage() should be removed
- [SENTRY-2019](#) - Improve logging in SentryPlugin
- [SENTRY-2040](#) - When getting Snapshots from HMS we need more logging around cases when a snapshot is not being received
- [SENTRY-2066](#) - DB name is not set for AlterTable
- [SENTRY-2068](#) - Disable HTTP TRACE method from the Sentry Web Server
- [SENTRY-2078](#) - Have sentry server print an obvious INFO level log message when it becomes the writer
- [SENTRY-2120](#) - cross-site scripting vulnerability in LogLevelServlet
- [SENTRY-2124](#) - LeaderStatusMonitor.toString() throws IllegalFormatException with AtomicLong
- [SOLR-11873](#) - HA cases fixed when NameNode fails over. Solr will remain green in CM and will remain fast due to remembering the last known good NameNode.
- [SPARK-16845](#) - [SQL][BRANCH-1.6] GeneratedClass\$SpecificOrdering` grows beyond 64 KB
- [SPARK-21551](#) - [PYTHON] Increase timeout for PythonRDD.serveIterator
- [SPARK-23438](#) - [DSTREAMS] Fix DStreams data loss with WAL when driver crashes
- [SQOOP-2976](#) - Introduce a flag to expand decimal values to fit AVRO schema
- [SQOOP-3233](#) - Fixed error in HCatalog imports that occurred when the target column type was either CHAR or VARCHAR and the source column was a number (decimal or int)
- [SQOOP-3241](#) - From now on, the ImportAllTables tool will use a clone SqoopOptions object for every import instead of using the same. This fix also includes tests to ensure that cloning works correctly
- [ZOOKEEPER-1872](#) - QuorumPeer is not shutdown in few cases
- [ZOOKEEPER-2249](#) - CRC check failed when preAllocSize smaller than node data

Issues Fixed in CDH 5.14.1

CDH 5.14.1 fixes the following issue:

Hive Replications Can Fail Intermittently

This issue, which was also reported in [TSB-2018-276](#), is fixed in CDH 5.14.1. For further details, see the original [Known Issue](#).

Issues Fixed in CDH 5.14.0

CDH 5.14.0 fixes the following issues.

Hue

Hue operations randomly fail and display a 502 error

Hue operations randomly fail.

Affected versions: CDH 5.12.x, 5.13.0

Workaround: Make configuration changes to the Hue Load Balancer:

1. Find the Hue Load Balancer host by selecting Host > All Host and search for "Hue Load Balancer"
2. Navigate to the following directory on the Hue Load Balancer host: `/usr/share/cmf/hue/httpd/`.
3. Edit the `httpd.conf` file and add the following property after the `ProxyPass` property: `SetEnv proxy-intial-not-pooled 1`.

4. Inside the `<IfDefine Cloudera_HTTPD_USE_SSL>` block under the `SSLProtocol` property line, add the following properties:

```
SSLProxyProtocol all -SSLv2 -SSLv3
SSLCipherSuite
```

Cloudera Bug: CDH-57588

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.14.0:

- [FLUME-2579](#) - Support durable subscription in JMSSource
- [FLUME-3046](#) - Kafka Sink and Source Configuration Improvements
- [FLUME-3100](#) - Support arbitrary header substitution for topic of Kafka Sink
- [FLUME-3158](#) - Upgrade surefire version and config
- [FLUME-3195](#) - Split the KafkaChannelTest to avoid timeouts
- [FLUME-3201](#) - Fix SyslogUtil to handle RFC3164 format in December correctly
- [HADOOP-10829](#) - Iteration on CredentialProviderFactory.serviceLoader is thread-unsafe
- [HADOOP-12717](#) - NPE when trying to rename a directory in Windows Azure Storage FileSystem
- [HADOOP-12903](#) - IPC Server should allow suppressing exception logging by type, not log 'server too busy' messages.
- [HADOOP-14475](#) - Metrics of S3A don't print out when enabled
- [HADOOP-14687](#) - AuthenticatedURL will reuse bad/expired session cookies
- [HADOOP-14814](#) - Fix incompatible API change on FsServerDefaults to HADOOP-14104
- [HADOOP-14864](#) - FSDataInputStream#unbuffer UOE should include stream class name
- [HADOOP-14880](#) - [KMS] Document&test missing KMS client side configs
- [HADOOP-14891](#) - Remove references to Guava Objects.toStringHelper
- [HADOOP-14892](#) - MetricsSystemImpl should use Time.monotonicNow for measuring durations
- [HADOOP-14919](#) - BZip2 drops records when reading data in splits
- [HADOOP-14944](#) - Add JvmMetrics to KMS.
- [HADOOP-14948](#) - Document missing config key hadoop.treat.subject.external
- [HADOOP-14960](#) - Add GC time percentage monitor/alerter
- [HADOOP-15023](#) - ValueQueue should also validate (int) (lowWatermark * numValues) > 0 on construction.
- [HADOOP-15067](#) - GC time percentage reported in JvmMetrics should be a gauge, not counter
- [HDFS-9941](#) - Do not log StandbyException on NN, other minor logging fixes
- [HDFS-10326](#) - Disable setting tcp socket send/receive buffers for write pipelines
- [HDFS-10984](#) - Expose nntop output as metrics
- [HDFS-11659](#) - TestDataNodeHotSwapVolumes.testRemoveVolumeBeingWritten fail due to no DataNode available for pipeline recovery
- [HDFS-11799](#) - Introduce a config to allow setting up write pipeline with fewer nodes than replication factor
- [HDFS-11896](#) - Non-dfsUsed will be doubled on dead node re-registration
- [HDFS-12052](#) - Set SWEBHDFS delegation token kind when ssl is enabled in HttpFS
- [HDFS-12143](#) - Improve performance of getting and removing inode features
- [HDFS-12171](#) - Reduce IIP object allocations for inode lookup
- [HDFS-12177](#) - NameNode exits due to setting BlockPlacementPolicy loglevel to Debug
- [HDFS-12319](#) - DirectoryScanner will throw IllegalStateException when Multiple BP's are present
- [HDFS-12371](#) - BlockVerificationFailures and BlocksVerified show up as 0 in Datanode JMX
- [HDFS-12485](#) - expunge may fail to remove trash from encryption zone
- [HDFS-12495](#) - TestPendingInvalidateBlock#testPendingDeleteUnknownBlocks fails intermittently
- [HDFS-12544](#) - SnapshotDiff - support diff generation on any snapshot root descendant directory.
- [HDFS-12578](#) - TestDeadDatanode#testNonDFSUsedONDeadNodeReReg failing in branch-2.7
- [HDFS-12603](#) - Enable async edit logging by default

- [HDFS-12754](#) - Lease renewal can hit a deadlock
- [MAPREDUCE-5124](#) - AM lacks flow control for task events
- [MAPREDUCE-6797](#) - Job history server scans can become blocked on a single, slow entry
- [MAPREDUCE-7028](#) - Concurrent task progress updates causing NPE in Application Master
- [YARN-2162](#) - Add ability in Fair Scheduler to optionally configure maxResources in terms of percentage.
- [YARN-7049](#) - FSAppAttempt preemption related fields have confusing names.
- [YARN-7057](#) - FSAppAttempt#getResourceUsage doesn't need to consider resources queued for preemption.
- [YARN-7363](#) - Revert "ContainerLocalizer don't have a valid log4j config in case of Linux container executor. (Contributed by Yufei Gu)"
- [YARN-7363](#) - ContainerLocalizer don't have a valid log4j config in case of Linux container executor.
- [HBASE-5291](#) - Add Kerberos HTTP SPNEGO authentication support to HBase web consoles
- [HBASE-15354](#) - Use same criteria for clearing meta cache for all operations
- [HBASE-15714](#) - We are calling checkRow() twice in doMiniBatchMutation()
- [HBASE-17424](#) - Disable external entity parsing in RemoteAdmin
- [HBASE-17590](#) - Drop cache hint should work on store file write path
- [HBASE-17658](#) - Fix bookkeeping error with max regions for a table
- [HBASE-18125](#) - shell disregards spaces at the end of a split key in a split file
- [HBASE-18141](#) - Regionserver fails to shutdown when abort triggered during RPC call
- [HBASE-18142](#) - Deletion of a cell deletes the previous versions too
- [HBASE-18233](#) - 'We shouldn't wait for readlock in doMiniBatchMutation in case of deadlock'
- [HBASE-18471](#) - The DeleteFamily cell is skipped when StoreScanner seeks to next column
- [HBASE-18480](#) - The cost of BaseLoadBalancer.cluster is changed even if the rollback is done
- [HBASE-18572](#) - Delete can't remove the cells which have no visibility label
- [HBASE-18665](#) - ReversedScannerCallable invokes getRegionLocations incorrectly
- [HBASE-18675](#) - Making {max,min}SessionTimeout configurable for MiniZooKeeperCluster
- [HBASE-18743](#) - HFiles in use by a table which has the same name and namespace with a default table cloned from snapshot may be deleted when that snapshot and default table are deleted
- [HBASE-18757](#) - Fix improper bitwise & in bucketcache offset calculation
- [HBASE-18847](#) - Remove unneeded synchronized block from HFileWriterV2 warning in branch-1.2
- [HBASE-18890](#) - Backport HBASE-14499 (Master coprocessors shutdown will not happen on master abort) to branch-1
- [HBASE-18907](#) - Methods missing rpc timeout parameter in HTable
- [HBASE-18924](#) - Backport HBASE-18568 (Correct metric of numRegions) to branch-1.2 and branch-1.3
- [HBASE-18998](#) - processor.getRowsToLock() always assumes there is some row being locked
- [HBASE-19215](#) - Incorrect exception handling on the client causes incorrect call timeouts and byte buffer allocations on the server
- [HIVE-9447](#) - Metastore: inefficient Oracle query for removing unused column descriptors when add/drop table/partition
- [HIVE-10807](#) - Invalidate basic stats for insert queries if autogather=false
- [HIVE-12828](#) - Update Spark version to 1.6
- [HIVE-14170](#) - Beeline IncrementalRows should buffer rows and incrementally re-calculate width if TableOutputFormat is used
- [HIVE-14303](#) - CommonJoinOperator.checkAndGenObject should return directly to avoid NPE if ExecReducer.close is called twice.
- [HIVE-15166](#) - Provide beeline option to set the jline history max size
- [HIVE-15519](#) - BitSet not computed properly for ColumnBuffer subset
- [HIVE-16007](#) - When the query does not compile, LogRunnable never stops
- [HIVE-16708](#) - Exception thrown while renewing a Delegation Token
- [HIVE-16854](#) - SparkClientFactory is locked too aggressively
- [HIVE-17022](#) - Add mode in lock debug statements
- [HIVE-17157](#) - Add InterfaceAudience and InterfaceStability annotations for ObjectInspector APIs

- [HIVE-17192](#) - Add InterfaceAudience and InterfaceStability annotations for Stats Collection APIs
- [HIVE-17202](#) - Add InterfaceAudience and InterfaceStability annotations for HMS Listener APIs
- [HIVE-17203](#) - Add InterfaceAudience and InterfaceStability annotations for HCat APIs
- [HIVE-17362](#) - The MAX_PREWARM_TIME should be configurable on HoS
- [HIVE-17640](#) - Comparison of date return null if time part is provided in string
- [HIVE-17743](#) - Add InterfaceAudience and InterfaceStability annotations for Thrift generated APIs
- [HIVE-17969](#) - Metastore to alter table in batches of partitions when renaming table
- [HIVE-18009](#) - Multiple lateral view query is slow on hive on spark
- [HIVE-18228](#) - Azure credential properties should be added to the HiveConf hidden list
- [HIVE-18248](#) - Clean up parameters that need not change at run time
- [HUE-2173](#) - [dashboard] Convert timeline widget to th nested facets
- [HUE-2173](#) - [dashboard] Only add term prefix when it is not empty
- [HUE-2173](#) - [dashboard] Allow min and max metrics for date types
- [HUE-2173](#) - [dashboard] Show sorting and facet limit when it makes sense
- [HUE-2173](#) - [dashboard] Port the range & up next text facet
- [HUE-2173](#) - [search] Avoid breaking search bar layout on normal resolutions
- [HUE-2173](#) - [dashboard] Port timeline to the nested facets
- [HUE-2173](#) - [dashboard] Do not error on sorting on text facets
- [HUE-2173](#) - [dashboard] Port the gradient map to the nested faceting
- [HUE-2173](#) - [dashboard] Longer search bar when analytics is supported
- [HUE-2173](#) - [dashboard] Reformat some part of the code
- [HUE-2173](#) - [dashboard] Port timeline and gradient map widgets to recursive nested facets
- [HUE-2173](#) - [dashboard] Pass to clean-up recursive facets template
- [HUE-2173](#) - [dashboard] Do not show sorting option on hit widget
- [HUE-2173](#) - [dashboard] Disable removing the last nested facet of a widget
- [HUE-2173](#) - [dashboard] Add observable subscribes to newly added facets
- [HUE-2173](#) - [dashboard] Hit widget compatibility with recursive facets
- [HUE-2173](#) - [dashboard] First pass of backend to support recursive facets
- [HUE-2173](#) - [dashboard] Refactor model to support N dimensions: UI part
- [HUE-2173](#) - [search] Harmonize indexer to use simpler field types
- [HUE-2173](#) - [dashboard] Support new facets in old dashboards
- [HUE-2173](#) - [dashboard] Mini skeleton for static dashboard
- [HUE-2173](#) - [search] Protect against empty nested facet buckets when nothing matches
- [HUE-2451](#) - [dashboard] Use SQL autocomplete for the query input
- [HUE-2451](#) - [dashboard] Perform search on enter in the query editor
- [HUE-2451](#) - [dashboard] Make top terms limit parameterizable
- [HUE-2451](#) - [dashboard] Round the left border of the query input
- [HUE-2451](#) - [dashboard] Don't use the Solr query autocomplete when engines other than Solr are used
- [HUE-2451](#) - [dashboard] Use prefix when fetching terms for value autocomplete
- [HUE-2451](#) - [dashboard] Use the terms API instead of samples when suggesting values in the query autocomplete
- [HUE-2451](#) - [dashboard] Enable placeholder for the simple ace editor component
- [HUE-2451](#) - [dashboard] Suggest fields, sample values and keywords in the query input
- [HUE-2451](#) - [dashboard] Use the simple Ace editor as the query input in the search bar
- [HUE-2451](#) - [dashboard] Use flex layout for the top search bar to adjust the width of the query input automatically
- [HUE-2451](#) - [dashboard] Don't handle partial values in the query parser
- [HUE-2451](#) - [dashboard] Add query autocomplete of keyword matching, values and boolean expressions
- [HUE-2451](#) - [dashboard] Initial generation setup of Solr query parser
- [HUE-2451](#) - [dashboard] Rename the Solr expression parser to Solr formula parser
- [HUE-2451](#) - [core] Split parser generation in the makefile
- [HUE-3277](#) - [assist] Use the same viewmodel for HDFS, S3 and ADLS entries

- [HUE-5480](#) - [autocomplete] Suggest S3 and ADLS paths in the autocomplete
- [HUE-5568](#) - [editor] Grid result not correctly aligned when browser zoom is not 100%
- [HUE-6106](#) - [autocomplete] Add / as an initial suggestion for empty file paths
- [HUE-6129](#) - [metastore] Show actual sample API errors in the panel instead of generic message
- [HUE-6159](#) - [autocomplete] Don't mark column alias references as columns belonging to a table
- [HUE-6169](#) - [editor] More obvious query log access when there is no "No error message"
- [HUE-6238](#) - [useradmin] ConcurrentUserSessionMiddleware - missing user
- [HUE-6238](#) - [useradmin] Option to disable concurrent user sessions
- [HUE-6256](#) - Revert "[presentation] Turn off the feature flag by default"
- [HUE-6555](#) - [editor] Remove the refresh button of the Saved Query tabs
- [HUE-6555](#) - [editor] Avoid listing only a newly saved query in the Saved Queries tab
- [HUE-6650](#) - [autocomplete] Merge popular group and order by column suggestions with existing column suggestions
- [HUE-6688](#) - [autocomplete] Switch to foreachVisible for the autocomplete results
- [HUE-6688](#) - [autocomplete] Make sure all running ajax requests are cancelled when closing the autocomplete
- [HUE-6784](#) - [oozie] Disable editor access also disable the browser
- [HUE-6799](#) - [editor] Use the last cursor position when opening from the query history
- [HUE-6799](#) - [editor] Save the cursor position with the snippet
- [HUE-6835](#) - [autocomplete] Update Impala built-in UDFs
- [HUE-6875](#) - [home] Introduce top header like all the other apps
- [HUE-6958](#) - [autocomplete] Fix issue with autocomplete after Hive GROUPING SETS
- [HUE-6958](#) - [autocomplete] Improve autocomplete of Hive windowing statements
- [HUE-6958](#) - [autocomplete] Add autocomplete for Hive MERGE statements
- [HUE-6958](#) - [autocomplete] Update the Hive TABLESAMPLE autocomplete
- [HUE-6958](#) - [autocomplete] Add INNER JOIN to the Hive suggestions
- [HUE-6958](#) - [autocomplete] Update autocomplete for the Hive LIMIT clause
- [HUE-6958](#) - [autocomplete] Add support for Hive ABORT statement
- [HUE-6958](#) - [autocomplete] Improve Hive SHOW TBLPROPERTIES
- [HUE-6958](#) - [autocomplete] Add autocomplete support for Hive SHOW VIEWS statements
- [HUE-6958](#) - [autocomplete] Add autocomplete support for Hive DROP TEMPORARY FUNCTION statements
- [HUE-6958](#) - [autocomplete] Update autocomplete for Hive ALTER TABLE statements
- [HUE-6958](#) - [autocomplete] Update the Hive reserved keyword lists
- [HUE-6958](#) - [autocomplete] Update autocomplete for Hive CREATE TABLE statements
- [HUE-6958](#) - [autocomplete] Add the DOUBLE PRECISION data type to Hive
- [HUE-6958](#) - [autocomplete] Add support for Hive ALTER DATABASE statements
- [HUE-6998](#) - [frontend] Improve configurability of Leaflet maps
- [HUE-7049](#) - [frontend] Generify the context popover and get rid of unused code
- [HUE-7058](#) - [search] Enable right panel for the search app in Hue 4
- [HUE-7059](#) - [search] Add list of fields to the right panel
- [HUE-7062](#) - [core] Fix all the 'ascii' codec can't decode from Solr function descriptions
- [HUE-7062](#) - [dashboard] Always suggest aggregate functions in the formula autocomplete
- [HUE-7062](#) - [dashboard] Convert arithmetic operators to their corresponding functions in the formula editor
- [HUE-7062](#) - [dashboard] Move formula editor cursor between the parentheses and trigger autocomplete when inserting functions
- [HUE-7062](#) - [dashboard] Suggest fields in the formula editor
- [HUE-7062](#) - [dashboard] Prevent line breaks when typing fast in the single line Ace editor
- [HUE-7062](#) - [dashboard] Only suggest aggregate functions when the formula editor is empty
- [HUE-7062](#) - [frontend] Move context popover and related ko components to ko_components
- [HUE-7062](#) - [frontend] Extract all ko components to separate files
- [HUE-7062](#) - [dashboard] Improve the size of the formula editor
- [HUE-7062](#) - [autocomplete] Link formula to the autocomplete

- [HUE-7062](#) - [dashboard] Suggest Solr functions in the dashboard formula editor
- [HUE-7062](#) - [dashboard] Add Solr function library to the autocomplete
- [HUE-7062](#) - [autocomplete] Extract commonly used autocomplete functionality to SqlUtils
- [HUE-7062](#) - [frontend] Add a generic Ace editor component that supports the Hue autocomplete dropdown
- [HUE-7062](#) - [search] Add backend logic for formula and remove old ops
- [HUE-7062](#) - [search] Add custom autocomplete parser for solr formulas
- [HUE-7078](#) - [editor] Fix console error on hover over old syntax error token
- [HUE-7078](#) - [editor] Improve handling of backticked values in the syntax checker
- [HUE-7078](#) - [editor] Don't mark custom variables as missing
- [HUE-7126](#) - [oozie] Improve Spark action to pick up the hive-site.xml easier
- [HUE-7140](#) - [editor] Add select variables
- [HUE-7140](#) - [editor] Add select variables
- [HUE-7186](#) - [editor] Support commenting variables
- [HUE-7186](#) - [editor] Support commenting variables
- [HUE-7190](#) - [fb] Pressing back in fb when path is not accessible does nothing
- [HUE-7215](#) - [indexer] Avoid duplicated column names
- [HUE-7217](#) - [dashboard] Avoid 500 on dashboard when search is blacklisted
- [HUE-7220](#) - [dashboard] Correctly reload the loaded states of widgets
- [HUE-7220](#) - [dashboard] Also support downloading facet result in csv and xls
- [HUE-7220](#) - [dashboard] Support downloading individual facet data
- [HUE-7234](#) - [editor] New variable should be empty
- [HUE-7236](#) - [metastore] Sample data shows the lock row functionality, but it does not work. Same for expand functionality
- [HUE-7239](#) - [editor] Use only one instance of the web workers for all snippets
- [HUE-7244](#) - [presentation] Add feedback effect to the global cancel operation
- [HUE-7244](#) - [presentation] Cancel a running presentation on exit
- [HUE-7244](#) - [presentation] Support cancelling the main execution
- [HUE-7245](#) - [search] Protect buckets when other filters can empty their values
- [HUE-7245](#) - [dashboard] Only propose dimension operations compatible with the data type
- [HUE-7247](#) - [editor] Show custom context menu when dragging views into the editor
- [HUE-7248](#) - [adls] Added client_id & tenant_id script handling
- [HUE-7248](#) - [adls] Add support for password script
- [HUE-7248](#) - [adls] Remove compress feature
- [HUE-7248](#) - [adls] fix test_export_result
- [HUE-7248](#) - [adls] fix test_export_result
- [HUE-7248](#) - [adls] fix export to file
- [HUE-7248](#) - [adls] filter filesystems for copy/save
- [HUE-7248](#) - [adls] fix upload hdfs root
- [HUE-7248](#) - [adls] fix path autocomplete
- [HUE-7248](#) - [adls] fix upload + error on upload cancel
- [HUE-7248](#) - [adls] Avoid cyclic dependency that prevents login
- [HUE-7248](#) - [adls] Avoid core libs import dependencies on apps
- [HUE-7248](#) - [adls] fix test test__get_fs & test__get_fs_pair
- [HUE-7248](#) - [adls] fix file upload over 30mb, remove logging auth token
- [HUE-7248](#) - [adls] fix assist hdfs error on root
- [HUE-7248](#) - [adls] Fix test_remove and splitpath()
- [HUE-7248](#) - [adls] Ability to compress / uncompress data in ADLS
- [HUE-7248](#) - [importer] Automatically set as external table based on ADLS data
- [HUE-7248](#) - [adls] Fix path of the hueversion symlink
- [HUE-7248](#) - [adls] Ability to browse data / content in ADLS

- [HUE-7248](#) - [assist] Add custom icon for ADLS
- [HUE-7250](#) - [assist] Improve performance when showing a table in assist
- [HUE-7251](#) - [editor] Don't select all and move cursor on undo after paste
- [HUE-7252](#) - [importer] CSV with quoted new lines in fields are not handled properly
- [HUE-7254](#) - [editor] Update the Schedule tab when saving a new query without requiring a page refresh
- [HUE-7256](#) - [fb] Improve UX for compress/extract as batch job
- [HUE-7257](#) - [autocomplete] Remove the SET options for Hive
- [HUE-7257](#) - [autocomplete] Add autocompletion of SET options for Hive and Impala
- [HUE-7261](#) - [editor] Prevent enormous chart legend labels to take over the chart
- [HUE-7263](#) - [editor] Resize handle off when scrolled down on editor and variables are present
- [HUE-7263](#) - [editor] Resize handle off when scrolled down on editor and variables are present
- [HUE-7265](#) - [oozie] Avoid deadlock on multiple json GET to retrieve a coordinator
- [HUE-7266](#) - [oozie] Do not change the URL when saving a coordinator of an integrated scheduler
- [HUE-7268](#) - [assist] Enable right click on the current selected database
- [HUE-7269](#) - [editor] Error popup when exploring table from assist and NavOpt on
- [HUE-7272](#) - [autocomplete] Cycle the autocomplete suggestions on up/down at the limits
- [HUE-7275](#) - [jb] Extract / Compress job has wrong timestamps and task names
- [HUE-7278](#) - [editor] Fix JS error on missing parseLocation for the syntax checker
- [HUE-7279](#) - [editor] Include column aliases in possible suggestions for the syntax checker
- [HUE-7280](#) - [editor] Mark the executed statement in the gutter when multiple statements are selected
- [HUE-7281](#) - [editor] Mark the correct lines for errors in multi query execution
- [HUE-7282](#) - [jb] 'int' object has no attribute 'lower' when click on certain jobs
- [HUE-7284](#) - [editor] Enable the syntax checker by default
- [HUE-7286](#) - [editor] Move sql identifier object to the beginning of the tooltip
- [HUE-7286](#) - [editor] Only show the main type for columns in the tooltip
- [HUE-7286](#) - [editor] Show column type in the hover tooltip if the details of the column is cached
- [HUE-7287](#) - [editor] Fix presentation mode list of fields height calculation
- [HUE-7288](#) - [editor] Don't check for risks when a statement has syntax errors
- [HUE-7289](#) - [core] Add article to the switch version message
- [HUE-7290](#) - [editor] Make it possible to clear ignored syntax errors
- [HUE-7291](#) - [oozie] Avoid JS error when clicking on file chooser on Schedule submit modal
- [HUE-7292](#) - [core] Avoid JS error on window resize when the user has visited the server logs page
- [HUE-7293](#) - [home] Avoid using the class divider for the breadcrumbs
- [HUE-7294](#) - [assist] Remove the nav search from assist
- [HUE-7295](#) - [frontend] Fix JS error on clearing the global search input with a selected result entry
- [HUE-7296](#) - [assist] Style the document context panel
- [HUE-7296](#) - [assist] Show a generic details view of a document when we can't display the contents
- [HUE-7296](#) - [assist] Use the document context popover in the left assist
- [HUE-7296](#) - [frontend] Add initial document context panel for global search
- [HUE-7296](#) - [frontend] Don't mark invisible space in the Ace highlight binding
- [HUE-7296](#) - [frontend] Add the UUID to returned documents in the global search result
- [HUE-7296](#) - [frontend] Improve ApiHelper fetchDocument function to optionally also return the contents
- [HUE-7297](#) - [doc] re.sub needs to be compatible with Python 2.6
- [HUE-7297](#) - [jb] Support middle click on the top 'Jobs' menu link
- [HUE-7297](#) - [doc] Document search highlighting should be case insensitive
- [HUE-7299](#) - [solr] Use the proper LISTALIAS API in Solr 6+ to list aliases
- [HUE-7299](#) - [assist] Keep the type filter when navigating to different directories
- [HUE-7299](#) - [assist] Include directories in the assist document filter
- [HUE-7300](#) - [scheduler] Simpler names for tasks of schedules
- [HUE-7300](#) - [jb] Hide more job properties when in mini browser mode

- [HUE-7303](#) - [editor] Avoid result grid size expansion on presentation mode
- [HUE-7304](#) - [editor] Fix scroll rendering problems of the grid results in presentation mode
- [HUE-7305](#) - [desktop] Revert "HUE-7192 [desktop] Hue Search is failing with DatabaseError"
- [HUE-7308](#) - [assist] Refresh files in assist and browser
- [HUE-7311](#) - [editor] Limit asterisk expansion to the tables in the select statement
- [HUE-7312](#) - [core] Fix scroll to column on sample popup
- [HUE-7313](#) - [jib] Fix JS error on schedule detail page
- [HUE-7314](#) - [home] Use svg icons for new document and folder in home
- [HUE-7314](#) - [frontend] Add a plus svg addon
- [HUE-7315](#) - Revert "[core] Change HBase icon for assist and app"
- [HUE-7315](#) - [core] Change HBase icon for assist and app
- [HUE-7316](#) - [jib] Skeleton of Impala Query integration into Job Browser
- [HUE-7321](#) - [frontend] Close the global search results on show in assist
- [HUE-7322](#) - [frontend] Re-open the global search results on focus if a query has been entered
- [HUE-7323](#) - [frontend] Enable drag and drop from the global search results
- [HUE-7327](#) - [impala] CTAS does not trigger an assist refresh
- [HUE-7329](#) - [editor] Don't warn when the syntax checker can't find a token
- [HUE-7330](#) - [assist] Improve filter for collection fields
- [HUE-7332](#) - [editor] Make the syntax checker aware of database changes
- [HUE-7333](#) - [search] Simplifies menu text and add autogrow to query boxes
- [HUE-7333](#) - [search] Improve top bar after adding the right panel assistant
- [HUE-7336](#) - [search] Add button to add column to the left
- [HUE-7337](#) - [search] Do not trigger several time the search() on dashboard load
- [HUE-7337](#) - [search] Use the correct field type for the possible metrics
- [HUE-7337](#) - [search] Counter widget should be allowed even if there is no number columns
- [HUE-7337](#) - [search] Only show one available dimension in the hit-widget
- [HUE-7339](#) - [search] Bar chart with dimension one and an analytics facet errors on click
- [HUE-7340](#) - [search] Solr 7 now supports min() max() on singlevalued strings
- [HUE-7340](#) - [search] Add new function types stddev() and variance()
- [HUE-7341](#) - [autocomplete] Fix issue where the first folder is ignored for HDFS suggestions
- [HUE-7342](#) - [editor] Make the syntax checker aware of CTE aliases
- [HUE-7343](#) - [assist] Exclude CTE aliases from the tables in the right assistant
- [HUE-7344](#) - [indexer] Avoid default 500 call to rdbms when not initialized
- [HUE-7345](#) - [indexer] Accept x01 char as field separator
- [HUE-7345](#) - [indexer] Index a Hive table into a Collection
- [HUE-7346](#) - [dashboard] Do not refresh the document panel on non save as
- [HUE-7347](#) - [importer] Proper error message when input table does not exist
- [HUE-7348](#) - [search] Do not error when loading a non analytics saved dashboard
- [HUE-7348](#) - [indexer] Properly use the table field delimiter when indexing them
- [HUE-7349](#) - [jib] New jobs are not inserted without a page refresh on an empty page
- [HUE-7350](#) - [assist] Right click on collections title should show a context menu
- [HUE-7351](#) - [assist] Always show the filter input in the assist panels
- [HUE-7352](#) - [assist] Add type to the collection right assistant
- [HUE-7353](#) - [search] Timeline chart shouldn't degrade the page performance
- [HUE-7354](#) - [frontend] Changed mini JB icon to fa-tasks
- [HUE-7354](#) - [frontend] Reverted the counter to orange, made it more round and changed icon for the mini JB
- [HUE-7354](#) - [frontend] The JB status counter should be more prominent
- [HUE-7355](#) - [home] Landing on Hue home with a wrong folder ID doesn't give the option of navigating away
- [HUE-7356](#) - [frontend] Fix js error when clearing search if a result is selected
- [HUE-7357](#) - [frontend] Remove the old global search

- [HUE-7358](#) - [frontend] Enable show in assist from the context popover for databases
- [HUE-7359](#) - [assist] Enable show in assist for documents in the global search results
- [HUE-7360](#) - [assist] Unify dashboard assistant look & feel with the SQL assistant
- [HUE-7361](#) - [assist] Dashboard assistant long field names shouldn't go to a new line
- [HUE-7362](#) - [assist] Switch to matching all text values in the sql search instead of at least one
- [HUE-7362](#) - [assist] Improve the facet handling for the right assistant filter
- [HUE-7362](#) - [frontend] Improve the autocomplete logic of the inline completer
- [HUE-7362](#) - [assist] Enable filter on column types in the left assist
- [HUE-7362](#) - [assist] Use the new inline autocomplete input for assist filters
- [HUE-7362](#) - [assist] Include free text parts in the query specification from the global search parser
- [HUE-7362](#) - [assist] Extract a generic inline autocomplete component
- [HUE-7362](#) - [assist] Extend the global search parser to also include the facet values in the parse results
- [HUE-7363](#) - [assist] Allow drag and drop of fields into a dashboard
- [HUE-7365](#) - [oozie] Simplify top right menu buttons on workflow editor
- [HUE-7366](#) - [oozie] Simplify top right menu buttons on schedule editor
- [HUE-7367](#) - [oozie] Simplify top right menu buttons on bundle editor
- [HUE-7368](#) - [frontend] Upgrade the Cloudera UI css to the latest version
- [HUE-7369](#) - [frontend] Migrate all remaining color variables in the LESS files to variables
- [HUE-7374](#) - [home] The expanded search box shouldn't go under the breadcrumbs
- [HUE-7376](#) - [indexes] Number of fields and sample records are not always shown
- [HUE-7377](#) - [frontend] Reset list of globally uploaded files at every new global upload
- [HUE-7378](#) - [fb] File Save As modal input is not styled correctly
- [HUE-7379](#) - [fb] Avoid usernames and groups to go to a new line
- [HUE-7380](#) - [dashboard] Adjust spacing of dimensions to allow drag and drop
- [HUE-7381](#) - [assist] Enable show in assist from the right assistant
- [HUE-7382](#) - [frontend] Use the default SQL interpreter as the source type for global search result entries
- [HUE-7383](#) - [core] Show Notebook as first non SQL editor when enabled
- [HUE-7384](#) - [importer] Add support for boolean values to Solr collections
- [HUE-7385](#) - [solr] Do not install Solr 7 examples with auto creation of fields
- [HUE-7386](#) - [search] Use proper types for some of the collection examples
- [HUE-7387](#) - [search] New ptypes from Solr 7
- [HUE-7388](#) - [core] Include schedule job counts into the global job count widget
- [HUE-7391](#) - [autocomplete] Limit the completer to the active statement only
- [HUE-7392](#) - [editor] Simplify top right menu buttons on the editor
- [HUE-7393](#) - [editor] dash is not recognized in variables
- [HUE-7394](#) - [core] Have dev mode on by default in development
- [HUE-7394](#) - [frontend] Use the dev flag instead of debug_mode to prevent caching of resources
- [HUE-7394](#) - [core] Flag for a dev mode
- [HUE-7395](#) - [jb] Support opening-up a query job id query
- [HUE-7395](#) - [editor] Add built-in link to the mini query job browser
- [HUE-7395](#) - [jb] Enable mini Job Browser for queries
- [HUE-7396](#) - [dashboard] Adding config properties to enable query builder and report
- [HUE-7396](#) - [dashboard] Re-organize available layouts
- [HUE-7397](#) - [editor] Adding headers to clipboard copy data
- [HUE-7397](#) - [editor] Add number of rows in clipboard copy to warn on possible truncation
- [HUE-7400](#) - [editor] Made close logs icon more visible and reliable
- [HUE-7400](#) - [editor] Made close logs icon bigger
- [HUE-7400](#) - [editor] Moved bigger icons to the specifics CSS classes
- [HUE-7400](#) - [editor] Slightly bigger icon size for the dashboard icon in menu
- [HUE-7400](#) - [editor] Slightly bigger show log toggling icons

- [HUE-7401](#) - [indexes] Add multi index select via SHIFT binding to list index page
- [HUE-7401](#) - [editor] Port quick save to dashboard to use the SolrClient
- [HUE-7403](#) - [core] Add 4.1 release notes
- [HUE-7405](#) - [assist] S3 document filter loses focus on search
- [HUE-7409](#) - [search] More subtle way of converting old types to new types
- [HUE-7409](#) - [search] ptypes should properly match any type
- [HUE-7409](#) - [search] The collection schem config was being ignored
- [HUE-7410](#) - [core] Support opening mini job browser links directly on the correct section
- [HUE-7411](#) - [frontend] Cannot share documents on touchscreen machines
- [HUE-7413](#) - [fb] parent directory should not be selectable
- [HUE-7415](#) - [editor] The editor textarea should retain its size on new query
- [HUE-7416](#) - [useradmin] Only display edited users on group change
- [HUE-7417](#) - [editor] Variables can override default values depending order
- [HUE-7418](#) - [hive] Table stats do not have numRows if the stats are not accurate
- [HUE-7419](#) - [core] Update SAML libraries Fix undefined failure_function for handling SAML response parsing exception
- [HUE-7419](#) - [core] Update SAML libraries add dependencies defusedxml(0.4.1), future(0.16.0), and pycryptodomex
- [HUE-7419](#) - [core] Update SAML libraries upgrade djangoaml2 (0.13.0 --> 0.16.4) and pysaml2
- [HUE-7420](#) - [jb] Hide empty Impala tabs
- [HUE-7420](#) - [jb] Fix test jobbrowser.tests.TestImpalaApi.test_apps
- [HUE-7420](#) - [jb] Incorrect number of Impala queries
- [HUE-7420](#) - [jb] Adding Impala text to the query page to prevent confusions with other SQL engines
- [HUE-7420](#) - [editor] Open correct mini browser tab depending on the type of query
- [HUE-7420](#) - [jb] Fix mini Impala jb expand + Styling
- [HUE-7420](#) - [jb] Fix TestImpalaApi.test_app
- [HUE-7420](#) - [jb] Remove white space from impala title
- [HUE-7420](#) - [jb] Queries: bring up bottom panel
- [HUE-7420](#) - [jb] Queries: bring up bottom panel
- [HUE-7420](#) - [jb] Fix impala link in editor log
- [HUE-7420](#) - [jb] Fix impala duration & progress
- [HUE-7420](#) - [jb] Fix impala summary & plan whitespace + perf
- [HUE-7420](#) - [jb] query browser 'NoneType' object has no attribute 'group'
- [HUE-7420](#) - [jb] queries are not in sorted order
- [HUE-7420](#) - [jb] total number of impala queries incorrect
- [HUE-7420](#) - [jb] query plan border color
- [HUE-7420](#) - [adls] replaced REFRESH_URL for TENANT_ID
- [HUE-7420](#) - [jb] enable query browser by default
- [HUE-7420](#) - [jb] impala kill job + filter
- [HUE-7420](#) - [jb] Impala Query integration
- [HUE-7424](#) - [frontend] Remove double jQuery library
- [HUE-7425](#) - [frontend] Add gridster.js library
- [HUE-7426](#) - [frontend] Add gridster Knockout binding
- [HUE-7428](#) - [core] Opening a view directly from top search lands on a 404
- [HUE-7431](#) - [core] When idle session timeout is enabled it causes issues with IDP initiated SAML
- [HUE-7432](#) - [frontend] Upgrade Knockout to the latest
- [HUE-7433](#) - [frontend] Add switch to enable new charting library
- [HUE-7434](#) - [core] Support more than one level URLs for custom apps in Hue 4
- [HUE-7435](#) - [frontend] Extract Leaflet renderers to a separate common file
- [HUE-7436](#) - [frontend] Avoid overriding of charting data transformers between editor and dashboard
- [HUE-7437](#) - [dashboard] History back sometimes shows the list of dashboards

- [HUE-7440](#) - [core] Remove load balancer cookie on login if not authenticated
- [HUE-7440](#) - [core] Remove load balancer cookie in logout view
- [HUE-7441](#) - [core] Allow to blacklist all the apps but dbms and file browser
- [HUE-7443](#) - [dashboard] Fix widget add dropdown menu style
- [HUE-7444](#) - [dashboard] Fix missing d3v3.tip declaration
- [HUE-7445](#) - [editor] Add session id to the Session panel
- [HUE-7447](#) - [core] Skip test on last access time as feature disabled
- [HUE-7447](#) - [core] Convert logged REST path to proper unicode
- [HUE-7447](#) - [core] Log metrics into debug info
- [HUE-7447](#) - [core] Add thread stack traces to debug log
- [HUE-7447](#) - [metadata] Add call tracking to navop calls
- [HUE-7447](#) - [core] Add thread tracebacks into the logs in debug mode
- [HUE-7447](#) - [core] Stop logging user pages counts
- [HUE-7447](#) - [core] Track return call of REST calls
- [HUE-7447](#) - [core] Log access page return time
- [HUE-7448](#) - [indexer] Support CTRL+A as field delimiter
- [HUE-7449](#) - [editor] Don't mark columns from CTEs as missing
- [HUE-7450](#) - [editor] Fix broken jasmine tests
- [HUE-7450](#) - [editor] Don't suggest aliases when a qualified identifier isn't found by the syntax checker
- [HUE-7458](#) - [editor] Fix js error in syntax checker for subquery references
- [HUE-7462](#) - [assist] Highlight matching parts of the functions filter
- [HUE-7463](#) - [editor] Don't mark Impala complex types as missing
- [HUE-7463](#) - [autocomplete] Improve location handling for map and array column references in Hive
- [HUE-7464](#) - [editor] Refactor SQL sessions to be less confusing
- [HUE-7464](#) - [editor] Do not save SQL sessions as they become stale
- [HUE-7464](#) - [editor] Inform if the session was recreated or not in the create_session API
- [HUE-7464](#) - [editor] Show log panel by default
- [HUE-7466](#) - [dashboard] Fix alignment of facet add typeahead
- [HUE-7467](#) - [frontend] Fix missing subtitle for the popovers
- [HUE-7467](#) - [frontend] Create Solr sample popup
- [HUE-7468](#) - [assist] Convert collection assist into tree mode
- [HUE-7471](#) - [assist] Don't show a scroll bar for entries with a spinner
- [HUE-7474](#) - [impala] Log the query planning phase
- [HUE-7475](#) - [core] Merge HUE-3102 to pysaml2 v4.4.0
- [HUE-7478](#) - [indexes] Re-use Solr sample popover in the index field list
- [HUE-7480](#) - [oozie] Also check for group permission for dashboard action
- [HUE-7481](#) - [frontend] Close the global search when clicking on show in assist or table browser
- [HUE-7482](#) - [metadata] Display last time of data refresh if available
- [HUE-7483](#) - [backend] validate property in hue.ini is ignored
- [HUE-7484](#) - [editor] Avoid old headers sticking for one second on result full screen mode exit
- [HUE-7485](#) - [frontend] Allow partial matching for facet values in the inline autocomplete
- [HUE-7486](#) - [oozie] 500 error when trying to submit Hive Workflow when beeswax blacklisted
- [HUE-7487](#) - [editor] Propose setting 'max_row_size' on error
- [HUE-7490](#) - [assist] File assist does not have an initial spinner
- [HUE-7491](#) - [fb] Newly uploaded file via the Upload button does not get the blue selection anymore
- [HUE-7492](#) - [fb] Second upload from Upload button seems to error
- [HUE-7493](#) - [editor] Remove fuzzy feeling from the autocompleter match
- [HUE-7494](#) - [core] RM HA does not seem to get persisted
- [HUE-7495](#) - [editor] Copy to clipboard adds an extra \t at the beginning of header line
- [HUE-7496](#) - [core] Bump cherryry threads by 10

- [HUE-7497](#) - [indexer] Fix the checking of an existing index or not while indexing
- [HUE-7497](#) - [indexer] Fix the link to an existing collection
- [HUE-7497](#) - [indexer] Refresh assist collection list on index creation
- [HUE-7498](#) - [importer] Also show file type on create index wizard from assist panel
- [HUE-7498](#) - [importer] Avoid js error when switching input types on a badly read data sample
- [HUE-7499](#) - [dashboard] Fix general typeahead positioning
- [HUE-7500](#) - [frontend] Upgrade Nicescroll to the latest
- [HUE-7501](#) - [dashboard] Aggregate field does not exist in older samples
- [HUE-7501](#) - [dashboard] Do not break the page when gridster is off
- [HUE-7501](#) - [dashboard] MVP of the gridster integration under a feature flag
- [HUE-7501](#) - [dashboard] Keep old layout intact and return a list of the new format
- [HUE-7501](#) - [dashboard] Convert layout to gridster format
- [HUE-7501](#) - [dashboard] Add configuration flag to enable the new layout
- [HUE-7502](#) - [aws] upload file to s3 location throws error: undefined
- [HUE-7503](#) - [dashboard] Rename generated formula to plain formula
- [HUE-7503](#) - [search] Keep sorting toggles in sync between facets of same dimension
- [HUE-7503](#) - [dashboard] Generate sort within the same dimension facet
- [HUE-7503](#) - [dashboard] Step 1 of sorting on aggregates of any dimension
- [HUE-7503](#) - [dashboard] Add sorting to Nth dimension
- [HUE-7503](#) - [dashboard] Add more than one dimension
- [HUE-7503](#) - [dashboard] Use new metric template for dimension 1
- [HUE-7503](#) - [dashboard] Analytics lean-up for counter widget
- [HUE-7503](#) - [dashboard] Move nested facet properties to measure template
- [HUE-7503](#) - [dashboard] Only show Add measure button when a field is selected
- [HUE-7503](#) - [dashboard] Convert add new measure to the new measure-form
- [HUE-7503](#) - [dashboard] New measure template
- [HUE-7503](#) - [dashboard] Simpler dimensions template
- [HUE-7504](#) - [editor] Prevent an infinite loop when changing the snippet type
- [HUE-7505](#) - [fb] Add a visual hint (green) to the upload status bar when successfully uploaded, and same for failures
- [HUE-7505](#) - [fb] Home link has a different highlight background than the path
- [HUE-7505](#) - [fb] The file properties in the file viewer are too close to the action menu
- [HUE-7505](#) - [fb] Style file editor should be the same as file viewer
- [HUE-7506](#) - [editor] Use the inline autocomplete in the result column search
- [HUE-7507](#) - [frontend] Use the inline autocomplete in the context popover column filter
- [HUE-7508](#) - [editor] Have the column filter visible by default in the context popover
- [HUE-7510](#) - [editor] Improve cursor position after dragging tables and columns to the editor
- [HUE-7511](#) - [fb] Fix replication factor for the currently selected file
- [HUE-7512](#) - [home] Select json on the import documents modal is weirdly big
- [HUE-7513](#) - [frontend] Top global search input is smaller than the query blue button and misaligned
- [HUE-7516](#) - [fb] . dir can be selected and allowed to copy into itself
- [HUE-7517](#) - [frontend] Properly close the global search results when opening a document
- [HUE-7518](#) - [editor] Remove bottom margin from the result field list filter
- [HUE-7523](#) - [assist] Enable drag and drop of documents in the assist panel
- [HUE-7524](#) - [oozie] Avoid parsing XML with resolve_entities
- [HUE-7528](#) - [editor] Show partition key details for columns in the contexts popover
- [HUE-7529](#) - [metastore] Show primary keys in the column list for a table
- [HUE-7532](#) - [home] Move create folder to the create document menu
- [HUE-7541](#) - [assist] Missing styles from the Hue 3 assist
- [HUE-7542](#) - [metadata] Only cache for a limited amount of time the list of cluster ids
- [HUE-7544](#) - [editor] Make identifier existence validator less eager when editing the identifier

- [HUE-7545](#) - [search] Reset query string in URL when it becomes empty
- [HUE-7546](#) - [dashboard] Automatically apply height to the Gridster widgets on Leaflet draw
- [HUE-7546](#) - [dashboard] Automatically apply height to the Gridster widgets on load and on plotly afterplot
- [HUE-7546](#) - [dashboard] Auto calculate height of tile based on content
- [HUE-7547](#) - [editor] Improve location handling for Impala complex identifiers
- [HUE-7548](#) - [metastore] Create table time conversion to local time can fail
- [HUE-7549](#) - [metastore] Drop table does not refresh the page
- [HUE-7551](#) - [autocomplete] Fix issue where suggestions are missing when there's no alias for a complex reference
- [HUE-7553](#) - [job] Attempt object has no attribute 'task_attempt_counters'
- [HUE-7554](#) - [job] Job Browser should provide YARN "Diagnostics" info when non empty
- [HUE-7555](#) - [importer] Type guessing for table sets everything as string
- [HUE-7555](#) - [core] Rename Amin section link in the top right user menu
- [HUE-7556](#) - [assist] Improve icons and i18n for the collections panel
- [HUE-7557](#) - [assist] Wire in the right sample popup for the collections items
- [HUE-7558](#) - [assist] Update the context menu for the collections items
- [HUE-7559](#) - [assist] Pre-select the source and root item for the collections panel
- [HUE-7560](#) - [assist] Added Solr specific facet values for the filter
- [HUE-7560](#) - [assist] Extracted filtered computed creation to AssistantUtils
- [HUE-7560](#) - [assist] Removed collection specific panel and re-used db one
- [HUE-7560](#) - [assist] Simplified loading of the right assist collection panel
- [HUE-7560](#) - [assist] Fix filtering of fields in the right assist collection panel
- [HUE-7560](#) - [assist] Re-use the new collections panel on the dashboard right assist
- [HUE-7561](#) - [sentry] Bump single user caching for metadata search
- [HUE-7562](#) - [core] Allow building Hue with Java 9
- [HUE-7563](#) - [metastore] Fields in nested column have a leading comma when preceded by array in type struct
- [HUE-7564](#) - Remove one of the get_connection(lldap_config) function in ldap_access.py.
- [HUE-7565](#) - [dashboard] Allow dragging to dashboard to create a new widget
- [HUE-7566](#) - [dashboard] Allow removing a Gridster widget
- [HUE-7567](#) - [all] Fix test desktop.lib.rest.test_clear_cookies
- [HUE-7568](#) - [frontend] Show description match in bold instead of italic for the top search results
- [HUE-7568](#) - [core] Top search document description can contains for matching highlights
- [HUE-7569](#) - [core] Adding Thrift arguments to the logging of the return call
- [HUE-7569](#) - [core] Finer grained Thrift pool instrumentation
- [HUE-7570](#) - [core] Bump ConnectionPooler to have a poolsize of CHERRYPY_SERVER_THREADS
- [HUE-7571](#) - [doc] Update thirdparty README
- [HUE-7572](#) - [editor] Show the complete statement in the document context popover
- [HUE-7573](#) - [assist] Sort the documents in the assist panel according to last modified time with folders on top
- [HUE-7574](#) - [assist] Create index from the Solr assist
- [HUE-7575](#) - [assist] Hide filter icon from the Solr assist
- [HUE-7576](#) - [assist] List collections aliases on the Solr assist
- [HUE-7577](#) - [editor] Don't assume databases have been loaded for all sources
- [HUE-7578](#) - [assist] Fetch more entries when scrolling to the end in the S3 and ADLS panels
- [HUE-7579](#) - [frontend] Fix anonymous user render of login page
- [HUE-7580](#) - [metadata] Split in two ddl and stats uploads
- [HUE-7583](#) - [editor] Export result to table throws already exists error
- [HUE-7584](#) - [solr] Support default field via the config API
- [HUE-7584](#) - [solr] Support deleting configSet via the config API
- [HUE-7584](#) - [solr] Support creating configSet via the config API
- [HUE-7585](#) - [indexer] Delete collection when data indexing fails in the wizard
- [HUE-7585](#) - [core] Tweak config creation to not override the base one

- [HUE-7585](#) - [core] Integrate Sentry protected schemas in config creation
- [HUE-7585](#) - [core] Properly detect if Solr is kerberized but not Sentry protected
- [HUE-7586](#) - [assist] Re-wire collection refresh subscriber
- [HUE-7587](#) - [indexer] Fix csv format type guessing test
- [HUE-7587](#) - [search] Properly create collection with default field when specified
- [HUE-7588](#) - [core] Option to disable Hue 3 switching
- [HUE-7591](#) - [pig] Also add hive to the sharelib when using hcatalog
- [HUE-7592](#) - [jb] Hide older browser if the new ones is enabled
- [HUE-7594](#) - [core] Adjust project license to be valid
- [HUE-7595](#) - [editor] The spacing of the new notebook icon is off
- [HUE-7596](#) - [assist] Avoid fetching the top columns for the collections panel
- [HUE-7597](#) - [core] Update 3rd party Readme with frontend libraries
- [HUE-7599](#) - [autocomplete] Remove some unsuitable SET options
- [HUE-7600](#) - [assist] Split the column entry assist for the right assistant
- [HUE-7601](#) - [indexer] Fix sample popover integration
- [HUE-7602](#) - [assist] Remove unused AssistCollectionEntry references
- [HUE-7603](#) - [dashboard] Smaller default empty widget
- [HUE-7603](#) - [dashboard] Initial layout chooser should create the layout in Gridster too
- [HUE-7610](#) - [dashboard] Add delete functionality to an empty widget
- [HUE-7611](#) - [dashboard] Add drop functionality to an empty widget
- [HUE-7612](#) - [dashboard] Recalculate height of non graphical widgets
- [HUE-7613](#) - [dashboard] Make resize icon more prominent
- [HUE-7614](#) - [dashboard] Enable drag'n'drop from assist
- [HUE-7615](#) - [doc] Add upgrade instructions
- [HUE-7616](#) - [editor] Shell job fails with TypeError: for +: 'dict' and 'str'
- [HUE-7617](#) - [fb] Compress can fail with output.properties data exceeds its limit
- [HUE-7619](#) - [indexer] Morphline indexer should not be on by default yet
- [HUE-7620](#) - [indexer] Offer to index split a field into a multi-valued field
- [HUE-7622](#) - [jb] Reruning a wokflow redirects to a json result
- [HUE-7624](#) - [core] Fix LDAP unit test
- [HUE-7624](#) - [core] Support multi-authentication with AllowFirstUserDjangoBackend and LdapBackend
- [HUE-7625](#) - [assist] Sort the folders in the documents panel and home by name
- [HUE-7626](#) - [editor] Fix issue where incorrect keywords were no longer marked by the syntax checker
- [HUE-7627](#) - [editor] Fix scroll to column after visiting metastore
- [HUE-7627](#) - [fb] Fail batch extraction job when extracted files already exist in output folder
- [HUE-7628](#) - [assist] Accept different kind of dashboard engines for the right assistant
- [HUE-7629](#) - [assist] Bump the field/column font size by one pixel
- [HUE-7630](#) - [dashboard] Add a basic query builder layout
- [HUE-7631](#) - [core] Keep older login URL for backward compatibility
- [HUE-7631](#) - Revert "[core] Have login page also prefixed with /hue""
- [HUE-7631](#) - Revert "[core] Have login page also prefixed with /hue"
- [HUE-7631](#) - [core] Have login page also prefixed with /hue
- [HUE-7632](#) - [dashboard] Set base size for each widget type to avoid expensive calculations
- [HUE-7633](#) - [dashboard] Resizing a Gridster widget should resize its content too
- [HUE-7633](#) - [dashboard] Resizing a Gridster widget should resize its content too
- [HUE-7634](#) - [dashboard] Persist and reload the gridster layout
- [HUE-7635](#) - [dashboard] Implement min count for SQL
- [HUE-7636](#) - [dashboard] Gridster should expand on assist collapse
- [HUE-7637](#) - [jb] Task attemps log size is very small
- [HUE-7642](#) - [editor] Anchor the marked syntax errors to have them move automatically with the text

- [HUE-7642](#) - [editor] Don't mark words as syntax errors while they're being typed
- [HUE-7643](#) - [editor] Only mark syntax errors for which there are suggestions
- [HUE-7643](#) - [editor] Prevent the syntax checker from marking stale editor identifiers as missing
- [HUE-7644](#) - [dashboard] Optimize initial search bar for smaller resolutions too
- [HUE-7647](#) - [core] Group sync fails when a non-standard (not uid) user_name_attr is used
- [HUE-7648](#) - [editor] Only mark missing identifiers when there are suggestions
- [HUE-7649](#) - [core] Add Jison to Notice file
- [HUE-7650](#) - [dashboard] Remove widget resize and move icons on Gridster mode
- [HUE-7651](#) - [dashboard] Show widget delete icon on hover on Gridster mode
- [HUE-7654](#) - [core] Little polishing of the new thread page
- [HUE-7654](#) - [core] Integrate Threads page as page on /about page
- [HUE-7655](#) - [dashboard] Add transitions to Gridster icons
- [HUE-7656](#) - [dashboard] Fix loading back the Gridster layout
- [HUE-7657](#) - [dashboard] Fix converting from old dashboard to Gridster
- [HUE-7658](#) - [dashboard] Improve widget margins
- [HUE-7660](#) - [editor] Update the Hive UDF library
- [HUE-7660](#) - [editor] Update the Impala UDF library
- [HUE-7661](#) - [frontend] Center the top search bar
- [HUE-7662](#) - [editor] Reloading a presentation mode page doesn't allow to load the editor again
- [HUE-7663](#) - [frontend] Add weights to global search result and keep selection after update
- [HUE-7663](#) - [frontend] Group all Hue apps in one global search category and only match if the name starts with the query
- [HUE-7663](#) - [frontend] Enable search for Hue apps in the global search
- [HUE-7665](#) - [editor] A failed query in presentation mode should reset the global play/stop icon
- [HUE-7666](#) - [editor] Presentation mode with variables shouldn't break the editor
- [HUE-7667](#) - [core] Fix http client pool by correctly mounting Transport Adapter to a prefix.
- [HUE-7668](#) - [frontend] Use up and down arrow keys to cycle suggestions in the inline autocomplete
- [HUE-7669](#) - [assist] Update the inline autocomplete on backspace
- [HUE-7670](#) - [frontend] Update the inline autocomplete after inserting suggestions
- [HUE-7671](#) - [assist] Fix the inline autocomplete in the dashboard assistant
- [HUE-7671](#) - [frontend] Include names in the inline autocomplete suggestions
- [HUE-7672](#) - [dashboard] The layout chooser should not error after creating a new dashboard programmatically
- [HUE-7673](#) - [dashboard] Enable vertical scroll inside the Gridster widgets
- [HUE-7676](#) - [core] skip_trash for hue_config_validation
- [HUE-7678](#) - [oozie] Running an old workflow from 3.7 might fail
- [HUE-7679](#) - [dashboard] Add clearable to the simple ace editor
- [HUE-7680](#) - Revert "[frontend] Get rid of Nicescroll to improve general performances"
- [HUE-7680](#) - Revert "[editor] Fix scroll in autocomplete dropdown"
- [HUE-7680](#) - Revert "[frontend] Tone down the scrollbars for webkit browsers"
- [HUE-7680](#) - Revert "[frontend] Improve CSS on not-so-top-of-the-line screens"
- [HUE-7680](#) - [frontend] Improve CSS on not-so-top-of-the-line screens
- [HUE-7680](#) - [frontend] Tone down the scrollbars for webkit browsers
- [HUE-7680](#) - [editor] Fix scroll in autocomplete dropdown
- [HUE-7680](#) - [frontend] Get rid of Nicescroll to improve general performances
- [HUE-7681](#) - [assist] Differentiate font sizes between the two assists
- [HUE-7682](#) - [assist] Expand the right assist content if there's just one table present
- [HUE-7683](#) - [assist] Add 'Open in assist' functionality for the right assist items
- [HUE-7684](#) - [assist] Add context menu to the collections title in left and right assist
- [HUE-7685](#) - [assist] Show table title in the right assist
- [HUE-7686](#) - [dashboard] Improve reliability of the db/table chooser

- [HUE-7687](#) - [metastore] Supply the source type when updating table and column comments
- [HUE-7688](#) - [metastore] Prevent the table browser from leaking css styles to other apps
- [HUE-7689](#) - [frontend] Avoid JS import error in old Hue
- [HUE-7692](#) - [editor] Upgrade failure from from CDH <= 5.4
- [HUE-7692](#) - [editor] Upgrade failure from from CDH <= 5.4
- [HUE-7693](#) - [assist] Make sure the right database is opened when a query is opened
- [HUE-7694](#) - [assist] Add sample tab to the collection context popover
- [HUE-7694](#) - [assist] Add show details to the collections name in left and right assist
- [HUE-7695](#) - [editor] Limit suggested values for the syntax checker
- [HUE-7695](#) - [editor] Only mark syntax errors in the active statement
- [HUE-7695](#) - [editor] Only identify complex locations for Impala
- [HUE-7695](#) - [editor] Limit background location checks to cached data only
- [HUE-7695](#) - [assist] Limit prefetching of columns in the right assist to the first 10 tables
- [HUE-7695](#) - [editor] Check identifier existence async
- [HUE-7697](#) - [core] Allow Kerberos Renewer to retry kinit multiple times
- [HUE-7697](#) - Revert "[core] Allow Kerberos Renewer to retry kinit multiple times"
- [HUE-7697](#) - [core] Allow Kerberos Renewer to retry kinit multiple times
- [HUE-7699](#) - [oozie] Shell document action are not listing any saved shell query
- [HUE-7701](#) - [search] Input focus reseted when filtering terms in column popup
- [HUE-7702](#) - [editor] Marker map lat and lon column section fields are always empty
- [HUE-7703](#) - [solr] Add collection sampling to the notebook API
- [HUE-7704](#) - [frontend] Fix wrong alignment of share dropdown options
- [HUE-7706](#) - [frontend] Fix js error in top search results
- [HUE-7707](#) - [assist] Show in assist should scroll just the left assist
- [HUE-7709](#) - [search] Dimension 2 facet selection display the same 1st column twice
- [HUE-7710](#) - [search] Prettify facet dimension humane description
- [HUE-7714](#) - [editor] Add column type to column name of row detail popup
- [HUE-7716](#) - [editor] Add a bit of left margin to the presentation mode
- [HUE-7717](#) - [editor] Double click on snippet title shouldn't add a new snippet in presentation mode
- [HUE-7718](#) - [editor] Enable presentation mode just for SQL editors
- [HUE-7719](#) - [editor] Consolidate the Pyspark mode
- [HUE-7720](#) - [home] Emptying the trash shouldn't show a constant spinner
- [HUE-7721](#) - [home] Improve "Empty trash" discoverability
- [HUE-7722](#) - [editor] Avoid JS error when user tries to execute zero queries in presentation mode
- [HUE-7723](#) - [editor] Only check for syntax in the Hive and Impala editors
- [HUE-7724](#) - [assist] Keep the state of the right assist tabs and refresh the assistant when needed
- [HUE-7725](#) - [autocomplete] Make sure we don't make suggestions on stale editor contents
- [HUE-7726](#) - [editor] Prevent statement parser error on '-' and '/'
- [HUE-7728](#) - [jib] 500% progress in live Impala query browser sometimes
- [HUE-7729](#) - [search] Support collections with hyphens in name
- [HUE-7730](#) - [fb] Refresh assist document list after creating workflow
- [HUE-7731](#) - [core] Move new and document list to top menu
- [HUE-7734](#) - [metadata] Avoid silent error when tenant id retrieval is empty
- [HUE-7735](#) - [editor] Prevent JS error from the risk suggestions
- [HUE-7736](#) - [assist] Make sure filtering the file panels resets the current page
- [HUE-7737](#) - [frontend] Make sure popovers are aligned when opened the first time
- [HUE-7739](#) - [metastore] Fix column context popover in the table browser
- [HUE-7744](#) - [dashboard] Enable Ace Hue theme in the single line Ace editor
- [HUE-7745](#) - [dashboard] Added square brackets to the supported operators
- [HUE-7745](#) - [dashboard] Create Solr keywords highlighter for Ace

- [HUE-7746](#) - [assist] Show a spinner when searching in the right assistant while columns are being loaded
- [HUE-7747](#) - [frontend] Show a spinner when the global search results are being updated
- [HUE-7748](#) - [frontend] Preselect the first entry in the global search result list
- [HUE-7749](#) - [editor] Don't mark columns from subqueries as missing
- [HUE-7752](#) - [core] Remove sample popup column filtering toggling icon
- [IMPALA-467](#) - enable disabled metric verification
- [IMPALA-992](#) - [DOCS] Document impala-shell 'rerun' command
- [IMPALA-1291](#) - Parquet read fails if io buffer size is less than the footer size
- [IMPALA-1575](#) - part 2: yield admission control resources
- [IMPALA-1575](#) - part 2: yield admission control resources
- [IMPALA-1575](#) - Part 1: eagerly release query exec resources
- [IMPALA-1767](#) - Adds predicate to test boolean values true, false, unknown.
- [IMPALA-2107](#) - [DOCS] Document base64*code() functions
- [IMPALA-2181](#) - Add query option levels for display
- [IMPALA-2190](#) - [DOCS] from_timestamp() and to_timestamp()
- [IMPALA-2234](#) - remove workaround from stress test
- [IMPALA-2235](#) - Fix current db when shell auto-reconnects
- [IMPALA-2250](#) - Make multiple COUNT(DISTINCT) message state workarounds
- [IMPALA-2281](#) - Replace FNV with FastHash in exchange nodes
- [IMPALA-2494](#) - Support for byte array encoded decimals in Parquet scanner
- [IMPALA-2636](#) - HS2 GetTables() returns TABLE_TYPE as TABLE for VIEW
- [IMPALA-2758](#) - Remove BufferedTupleStream::GetRows
- [IMPALA-2810](#) - Remove column stats restoration when altering table
- [IMPALA-3200](#) - [DOCS] Document user-facing aspects of new buffer pool
- [IMPALA-3360](#) - Codegen inserting into runtime filters
- [IMPALA-3437](#) - DECIMAL_V2: avoid implicit decimal->double conversion
- [IMPALA-3504](#) - [DOCS] Document utc_timestamp()
- [IMPALA-3516](#) - Avoid writing to /tmp in testing
- [IMPALA-3548](#) - Prune runtime filters based on query options in the FE
- [IMPALA-3613](#) - Avoid topic updates to unregistered subscriber instances
- [IMPALA-3642](#) - Adding backend addresses to error statuses for some scratch failures.
- [IMPALA-3877](#) - support unpatched LLVM
- [IMPALA-3897](#) - Codegen null-aware constant in PHJ::ProcessBuildBatch()
- [IMPALA-3998](#) - deprecate --refresh_after_connect
- [IMPALA-4082](#) - Remove todo item in getRegionsInRange
- [IMPALA-4177](#) - IMPALA-6039: batched bit reading and rle decoding
- [IMPALA-4236](#) - Codegen CopyRows() for select nodes
- [IMPALA-4252](#) - Min-max runtime filters for Kudu
- [IMPALA-4252](#) - Move runtime filters to ScanNode
- [IMPALA-4513](#) - Promote integer types for ABS()
- [IMPALA-4524](#) - Batch ALTER TABLE...ADD PARTITION calls.
- [IMPALA-4591](#) - Bound Kudu client error mem usage
- [IMPALA-4620](#) - Refactor evalcost computation in query analysis
- [IMPALA-4622](#) - [DOCS] New Kudu ALTER TABLE syntax
- [IMPALA-4623](#) - [DOCS] Document file handle caching
- [IMPALA-4669](#) - [KRPC] Add kudu_rpc library to build
- [IMPALA-4670](#) - Introduces RpcMgr class
- [IMPALA-4704](#) - Turns on client connections when local catalog initialized.
- [IMPALA-4736](#) - Add SIGUSR1 behavior to help string for 'minidump_path' flag
- [IMPALA-4786](#) - Clean up how ImpalaServers are created

- [IMPALA-4847](#) - Simplify HdfsTable block metadata loading code
- [IMPALA-4856](#) - Port data stream service to KRPC
- [IMPALA-4856](#) - IMPALA-4872: Include KRPC services in plan fragment's destinations
- [IMPALA-4856](#) - Rename thrift-deps to gen-deps
- [IMPALA-4863](#) - IMPALA-5311: Correctly account the file type and compression codec
- [IMPALA-4918](#) - Support getting column comments via HS2
- [IMPALA-4927](#) - Impala should handle invalid input from Sentry
- [IMPALA-4939](#) - IMPALA-4940: Decimal V2 multiplication
- [IMPALA-4964](#) - Fix Decimal modulo overflow
- [IMPALA-4985](#) - use parquet stats of nested types for dynamic pruning
- [IMPALA-5018](#) - Error on decimal modulo or divide by zero
- [IMPALA-5019](#) - Decimal V2 addition
- [IMPALA-5129](#) - Use Kudu's Kinit code to avoid expensive fork
- [IMPALA-5142](#) - EventSequence displays negative elapsed time.
- [IMPALA-5174](#) - Suppress kudu flags that aren't relevant to Impala
- [IMPALA-5199](#) - prevent hang on empty row batch exchange
- [IMPALA-5210](#) - Count rows and collection items in parquet scanner separately
- [IMPALA-5211](#) - Simplifying nullif conditional.
- [IMPALA-5211](#) - Simplifying ifnull/isnull/nvl where conditional is a literal.
- [IMPALA-5243](#) - Speed up code gen for wide Avro tables.
- [IMPALA-5250](#) - Unify decompressor output_length semantics
- [IMPALA-5307](#) - Part 4: copy out uncompressed text and seq
- [IMPALA-5307](#) - Part 2: copy out strings in uncompressed Avro
- [IMPALA-5307](#) - part 1: don't transfer disk I/O buffers out of parquet
- [IMPALA-5307](#) - Part 3: remove TODO from RCFile
- [IMPALA-5317](#) - add DATE_TRUNC() function
- [IMPALA-5341](#) - Avoid unintended filter out in fe test
- [IMPALA-5376](#) - Implement all TPCDS test cases or alternates for Impala.
- [IMPALA-5383](#) - [DOCS] Document unpartitioned Kudu tables
- [IMPALA-5394](#) - Change ThriftServer() to always use TAcceptQueueServer
- [IMPALA-5416](#) - Fix an impala-shell command recursion bug
- [IMPALA-5417](#) - make I/O buffer queue fixed-size
- [IMPALA-5425](#) - Add test for validating input when setting query options
- [IMPALA-5429](#) - Multi threaded block metadata loading
- [IMPALA-5448](#) - fix invalid number of splits reported in Parquet scan node
- [IMPALA-5501](#) - [DOCS] Clarify 'binaries' w.r.t. installation procedure
- [IMPALA-5525](#) - Extend TestScannersFuzzing to test uncompressed parquet
- [IMPALA-5529](#) - [DOCS] New trunc() signatures
- [IMPALA-5538](#) - Revert "Use explicit catalog versions for deleted objects"
- [IMPALA-5538](#) - Use explicit catalog versions for deleted objects
- [IMPALA-5541](#) - Reject BATCH_SIZE greater than 65536
- [IMPALA-5589](#) - Re-apply:change "set" in impala-shell to show empty string for unset query options
- [IMPALA-5589](#) - Revert "change "set" in impala-shell to show empty string for unset query options"
- [IMPALA-5589](#) - change "set" in impala-shell to show empty string for unset query options
- [IMPALA-5599](#) - Clean up references to TimestampValue in be/src.
- [IMPALA-5599](#) - Fix for mis-use of TimestampValue
- [IMPALA-5624](#) - Replace "ls -l" with opendir() in ProcessStateInfo
- [IMPALA-5638](#) - [DOCS] Add known issue for Impala-Kudu-Sentry issue
- [IMPALA-5653](#) - Remove "unlimited" process mem_limit option
- [IMPALA-5664](#) - Unix time to timestamp conversions may crash Impala

- [IMPALA-5668](#) - Fix cast(X as timestamp) for negative subsecond Decimals
- [IMPALA-5736](#) - Add impala-shell argument to set default query options
- [IMPALA-5789](#) - Add always_false flag in bloom filter
- [IMPALA-5836](#) - Improvements to Eclipse frontend configuration.
- [IMPALA-5844](#) - use a MemPool for expr result allocations
- [IMPALA-5846](#) - Fix output path for kudu libraries
- [IMPALA-5854](#) - Revert "Update external hadoop versions"
- [IMPALA-5854](#) - Update external hadoop versions
- [IMPALA-5860](#) - upgrade to LLVM 3.9.1
- [IMPALA-5870](#) - Improve runtime profile for partial sort
- [IMPALA-5873](#) - Check for existence of sync_file_range()
- [IMPALA-5894](#) - [DOCS] Clarify placement of STRAIGHT_JOIN hint
- [IMPALA-5895](#) - clean up runtime profile lifecycle
- [IMPALA-5902](#) - add ThreadSanitizer build
- [IMPALA-5905](#) - build-all-flag-combinations addendum
- [IMPALA-5905](#) - add script for all-build-options job
- [IMPALA-5908](#) - Allow SET to unset modified query options.
- [IMPALA-5912](#) - fix crash in trunc(..., "WW") in release build
- [IMPALA-5920](#) - addendum - add missing RAT check
- [IMPALA-5920](#) - Remove admission control dependency on YARN RM jar
- [IMPALA-5926](#) - Avoid printing expensive stack when closing a session
- [IMPALA-5927](#) - Fix enable_distcc for zsh
- [IMPALA-5932](#) - Improve transitive closure computation performance in FE
- [IMPALA-5940](#) - Avoid log spew by using Status::Expected()
- [IMPALA-5940](#) - Avoid log spew by using Status::Expected.
- [IMPALA-5941](#) - Fix Metastore schema creation in create-test-configuration.sh
- [IMPALA-5957](#) - print memory address, not memory
- [IMPALA-5965](#) - avoid per-value switch on NeedsConversionInline() in parquet
- [IMPALA-5966](#) - Fix the result file location of PlannerTest
- [IMPALA-5975](#) - Work around broken beeline clients
- [IMPALA-5976](#) - Remove equivalence class computation in FE
- [IMPALA-5986](#) - Correct set-option logic to recognize digits in names.
- [IMPALA-5987](#) - LZ4 Codec silently produces bogus compressed data for large inputs
- [IMPALA-5988](#) - optimise MemPool::TryAllocate()
- [IMPALA-5999](#) - Fix LLVM linkage errors due LibCache sync issues
- [IMPALA-6001](#) - Part 1: Increase log level during catalog update processing
- [IMPALA-6002](#) - Add a LLVM diagnostic handler for LLVM linker errors
- [IMPALA-6004](#) - Fix test_row_filters failure on ASAN
- [IMPALA-6009](#) - Upgrade Guava to 14.0.1
- [IMPALA-6012](#) - workaround - downgrade hive temporarily
- [IMPALA-6016](#) - Fix logging in TableLoadingMgr class
- [IMPALA-6019](#) - Remove dead code parallel-executor*
- [IMPALA-6021](#) - Revert "IMPALA-6009: Upgrade Guava to 14.0.1"
- [IMPALA-6027](#) - Retry downloading toolchain components.
- [IMPALA-6030](#) - Don't start coordinator specific thread pools if a node isn't a coordinator node
- [IMPALA-6040](#) - skip test_multi_compression_types where Hive isn't supported
- [IMPALA-6053](#) - Fix exception when storadelds don't match hosts
- [IMPALA-6054](#) - Parquet dictionary pages should be freed on dictionary construction
- [IMPALA-6055](#) - Fix hdfs encryption test far Hadoop 2.8+
- [IMPALA-6060](#) - Check the return value of JNI exception handling functions

- [IMPALA-6067](#) - Enable S3 access via IAM roles for EC2 VMs
- [IMPALA-6068](#) - Revert "Fix dataload for complextypes_fileformat"
- [IMPALA-6068](#) - Fix dataload for complextypes_fileformat
- [IMPALA-6069](#) - Fix CodegenAnyVal's handling of 'nan'
- [IMPALA-6076](#) - Parquet BIT_PACKED deprecation warning
- [IMPALA-6080](#) - clean up table descriptor handling
- [IMPALA-6081](#) - Fix test_basic_filters runtime profile failure
- [IMPALA-6084](#) - Avoid using of global namespace for llvm
- [IMPALA-6099](#) - Fix crash in CheckForAlwaysFalse()
- [IMPALA-6100](#) - increase test_exchange_delays timeout on slow builds
- [IMPALA-6105](#) - Clarify argument order in single_node_perf_run
- [IMPALA-6106](#) - handle comments before set in test parser
- [IMPALA-6108](#) - IMPALA-6070: Parallel data load (re-instated).
- [IMPALA-6108](#) - Revert "IMPALA-6070: Parallel data load."
- [IMPALA-6109](#) - xfail TestHdfsUnknownErrors::test_hdfs_safe_mode_error_255
- [IMPALA-6118](#) - Fix assertion failure in coordinator bloom filter updating
- [IMPALA-6121](#) - remove I/O mgr request context cache
- [IMPALA-6123](#) - Fix column order of a query test in test_inline_view_limit
- [IMPALA-6124](#) - Fix alter table ddl updates and test
- [IMPALA-6126](#) - ASAN detects heap-use-after-free in thrift-server-test
- [IMPALA-6127](#) - Fix timeout in TestRuntimeFilter.test_wait_time
- [IMPALA-6128](#) - Spill-to-disk Encryption(AES-CFB + SHA256) is slow
- [IMPALA-6134](#) - Update code base to use impala::ConditionVariable
- [IMPALA-6136](#) - Part 1: Query duration should not be normally negative.
- [IMPALA-6137](#) - fix text scanner split delim mem mgmt
- [IMPALA-6144](#) - PublishFilter() continues to run after query failure/cancellation
- [IMPALA-6148](#) - Specifying thirdparty deps as URLs
- [IMPALA-6151](#) - add query-level fragment/backend counters
- [IMPALA-6155](#) - Allow tests to pass when ORDER BY does not cover the query.
- [IMPALA-6164](#) - Fix stale query profile in TestAlwaysFalseFilter
- [IMPALA-6171](#) - Revert "IMPALA-1575: part 2: yield admission control resources"
- [IMPALA-6173](#) - Fix SHOW CREATE TABLE for unpartitioned Kudu tables
- [IMPALA-6183](#) - Fix Decimal to Double conversion
- [IMPALA-6184](#) - Clean up after ScalarExprEvaluator::Clone() fails
- [IMPALA-6187](#) - Fix missing conjuncts evaluation with empty projection
- [IMPALA-6188](#) - make test_top_n_reclaim less flaky
- [IMPALA-6198](#) - marks a test as debug-only
- [IMPALA-6201](#) - Fix test_basic_filters on ASAN
- [IMPALA-6206](#) - Fix data load failure with -notests
- [IMPALA-6217](#) - fix DCHECK in Parquet fuzz test
- [IMPALA-6220](#) - Revert "IMPALA-6128: Spill-to-disk Encryption(AES-CFB + SHA256) is slow"
- [IMPALA-6225](#) - Part 1: Query profile date-time strings should have ns precision.
- [IMPALA-6232](#) - Disable file handle cache by default
- [IMPALA-6241](#) - timeout in admission control test under ASAN
- [IMPALA-6256](#) - Incorrect principal will be used for internal connections if FLAGS_be_principal is set
- [IMPALA-6262](#) - Always initialize runtime profile for DataSink
- [IMPALA-6280](#) - Materialize TupleIsNullPredicate for insert sorts
- [IMPALA-6281](#) - Fix use-after-free in InitAuth()
- [IMPALA-6285](#) - Don't print stack trace on RPC errors.
- [IMPALA-6286](#) - Remove invalid runtime filter targets.

- [IMPALA-6291](#) - disable AVX512 codegen in LLVM
- [IMPALA-6292](#) - addendum: Fix failing test
- [IMPALA-6292](#) - Fix incorrect DCHECK in decimal subtraction
- [IMPALA-6298](#) - Skip test_profile_fragment_instances on local filesystem
- [IMPALA-6308](#) - Fix bad Status() usage in data-stream-sender.cc
- [IMPALA-5307](#) - ScannerContext API change
- [IMPALA-5493](#) - Add Protobuf to CMakeLists.txt
- [KUDU-1078](#) - Fixed an error message commonly found in tablet server logs indicating that operations were being read "from the future".
- [KUDU-1411](#) - HybridTime timestamp propagation now works in the Java client when using scan tokens.
- [KUDU-2044](#) - Tombstoned tablets no longer report metrics.
- [KUDU-2173](#) - Fixed a bug in the C++ client which could cause tablets to be erroneously pruned, or skipped, during certain scans, resulting in fewer results than expected being returned from queries. The bug only affected tables whose range partition columns are a proper prefix of the primary key.
- [KUDU-2188](#) - Published Kudu Java artifacts are now fully compatible with JRE 7 and JRE 8. There was previously a bug in the release process which made them compatible only with JRE 8.
- Fixed a typo in the list of default TLS ciphers used by Kudu servers. As a result, two additional cipher suites are now available:
 - ECDHE-RSA-AES128-SHA256 TLSv1.2 Kx=ECDH Au=RSA Enc=AES(128) Mac=SHA256
 - AES256-GCM-SHA384 TLSv1.2 Kx=RSA Au=RSA Enc=AESGCM(256) Mac=AEAD
- [KUDU-2231](#) - Sparse column predicates no longer cause excessive data-block reads.
- [OOZIE-1401](#) - PurgeCommand should purge the workflow jobs w/o end_time
- [OOZIE-3072](#) - oozie.service.HadoopAccessorService.action.configurations should overwrite default values set in Hadoop's configuration files
- [OOZIE-3078](#) - PasswordMasker throws NPE with null arguments
- [OOZIE-3103](#) - Improve log message regarding max concurrency for coordinators
- [OOZIE-3112](#) - SparkConfigurationService overwrites properties provided via --properties-file option in SparkAction
- [OOZIE-3132](#) - Instrument SLACalculatorMemory
- [PIG-3299](#) - Provide support for LazyOutputFormat to avoid creating empty files
- [SENTRY-1272](#) - Enable ALTERVIEW_RENAME and ALTERVIEW_AS operation in hive binding
- [SENTRY-1712](#) - Add trigger mechanism for Sentry to push full path snapshot to Name Node.
- [SENTRY-1906](#) - Sentry clients to retry connections to server with delay to avoid failing fast
- [SENTRY-1963](#) - Sentry JSON reporter should use regular implementation for local file system
- [SENTRY-1992](#) - Improve parameter handling for SentryGenericProviderBackend
- [SENTRY-2014](#) - incorrect handling of HDFS paths with multiple forward slashes
- [SENTRY-2031](#) - Add trigger mechanism for Sentry to pull full path snapshot from HMS
- [SENTRY-2032](#) - Leading Slashes need to be removed when creating HMS path entries
- [SENTRY-2035](#) - Metrics should move to destination atomically
- [SENTRY-2036](#) - sentry_sync_notifications() should set ID when it returns errors
- [SENTRY-2046](#) - Create a full snapshot if AUTHZ_PATHS_SNAPSHOT_ID is empty, even if HMS and Sentry Notifications are in sync
- [SENTRY-2047](#) - isEmptyTableCore method in SentryStore has references to MAuthzPathsMapping when it should be generic
- [SPARK-15273](#) - YarnSparkHadoopUtil#getOutOfMemoryErrorArgument should respect OnOutOfMemoryError parameter given by user
- [SQOOP-3196](#) - Modify MySQLAuthTest to use configurable test database parameters
- [SQOOP-3197](#) - Database credentials cannot be configured in PostgreSQL third party tests
- [SQOOP-3229](#) - Document how to run third party tests manually with databases running in docker
- [SQOOP-3243](#) - Importing BLOB data causes "Stream closed" error on encrypted HDFS
- [SQOOP-3257](#) - Sqoop does not log and print metastore password

- [ZOOKEEPER-1583](#) - Document maxClientCnxns in conf/zoo_sample.cfg
- [ZOOKEEPER-1584](#) - Adding mvn-install target for deploying the zookeeper artifacts to .m2 repository
- [ZOOKEEPER-2887](#) - define dependency versions in build.xml to be easily overridden in build.properties

Issues Fixed in CDH 5.13.x

The following topics describe issues fixed in CDH 5.13.x, from newest to oldest release. You can also review [What's New in CDH 5.13.x](#) on page 1817 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.13.3

Missing results in Hive, Spark, Pig, Custom MapReduce jobs, and other Java applications when filtering Parquet data written by Impala

This issue, which was also reported in [TSB-2018-300](#), is fixed in CDH 5.13.3 and higher. For further details, see the original [Known Issue for Hive](#), [Known Issue for Spark](#), [Known Issue for Pig](#), or the [Known Issue for MapReduce](#).

Slow or stalled queries under highly concurrent write workloads when Sentry + Hive are used, caused by a Hive metastore deadlock

This issue, which was also reported in [TSB-2018-294](#), is fixed in CDH 5.13.3. For further details, see the original [Known Issue](#).

Crash due to CHECK failure during compactions: `pv_delete_redo != nullptr`

This issue, which was also reported in [TSB-2018-298](#), is fixed in CDH 5.13.3 and 5.14.2. For further details, see the original [Known Issue](#).

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.13.3:

- [HADOOP-6852](#) - apparent bug in concatenated-bzip2 support (decoding)
- [HADOOP-12886](#) - Exclude weak ciphers in SSLFactory through ssl-server.xml
- [HDFS-6054](#) - MiniQJMHACluster should not use static port to avoid binding failure in unit test.
- [HDFS-10903](#) - Replace config key literal strings with config key names II: hadoop hdfs
- [HDFS-12386](#) - Add fserver defaults call to WebhdfsFileSystem.
- [HDFS-13040](#) - Kerberized inotify client fails despite kinit properly
- [HDFS-13115](#) - In getNumUnderConstructionBlocks(), ignore the inodelds for which the inodes have been deleted
- [HDFS-13170](#) - Port webhdfs unmaskedpermission parameter to HTTPFS
- [MAPREDUCE-7048](#) - Uber AM can crash due to unknown task in statusUpdate
- [MAPREDUCE-7053](#) - Timed out tasks can fail to produce thread dump
- [YARN-2416](#) - InvalidStateTransitionException in ResourceManager if AMLauncher does not receive response for startContainers() call in time
- [HBASE-15714](#) - We are calling checkRow() twice in doMiniBatchMutation()
- [HBASE-15866](#) - Split hbase.rpc.timeout into *.read.timeout and *.write.timeout
- [HBASE-17319](#) - Truncate table with preserve after split may cause truncation to fail
- [HBASE-17924](#) - Consider sorting the row order when processing multi() ops before taking rowlocks
- [HBASE-18233](#) - We shouldn't wait for readlock in
- [HBASE-19163](#) - "Maximum lock count exceeded" from region server's batch processing
- [HBASE-19188](#) - Build fails on branch-1 using maven-3.5.2
- [HBASE-19364](#) - Truncate_preserve fails with table when replica region > 1
- [HBASE-19391](#) - Calling HRegion#initializeRegionInternals from a region replica can still re-create a region directory
- [HBASE-19553](#) - Old replica regions should be cleared from AM memory after primary region split or merge
- [HBASE-19728](#) - Add lock to filesCompacting in all place.
- [HBASE-19934](#) - HBaseSnapshotException when read replicas is enabled and online snapshot is taken after region splitting
- [HBASE-20052](#) - TestRegionOpen#testNonExistentRegionReplica fails due to NPE

- [HBASE-20061](#) - HStore synchronized member variable filesCompacting should be private
- [HIVE-14792](#) - AvroSerde reads the remote schema-file at least once per mapper, per table reference.
- [HIVE-16886](#) - HMS log notifications may have duplicated event IDs if multiple HMSs are running concurrently
- [HIVE-18671](#) - Lock not released after Hive on Spark query was cancelled
- [HIVE-18917](#) - Add spark.home to hive.conf.restricted.list
- [HUE-7431](#) - [core] When idle session timeout is enabled it causes issues with IDP initiated SAML
- [HUE-7920](#) - [frontend] The stacked graph shouldn't have empty areas in the stacks
- [HUE-7932](#) - [oozie] Clicking on workflow with several decision nodes hangs Hue
- [HUE-7992](#) - [oozie] Big workflows show horizontally scroll in Hue 4
- [HUE-7997](#) - [backend] Python LDAP library exposes password
- [HUE-8009](#) - [useradmin] Fix login page error for invalid local user when LDAP also enabled
- [HUE-8024](#) - [indexer] Use old collection API to avoid grant Update privilege
- [HUE-8032](#) - [core] XSS in Hue Admin Interface
- [HUE-8053](#) - [useradmin] LDAP authentication with sync_groups_on_login=true fails with KeyError exception
- [HUE-8055](#) - [desktop] Support multiple LDAP servers in LDAP Test command
- [HUE-8063](#) - [fb] Better error message when HTTPFS is not working
- [HUE-8069](#) - [core] Pyformance reporter thread stops on exceptions in guage callbacks
- [HUE-8069](#) - [core] Fix code format and exception logging
- [IMPALA-6353](#) - Fixed the crash in snappy decompressor.
- [IMPALA-6473](#) - Error in analytic sort with same expr in 'partition by' and 'order by'.
- [IMPALA-6500](#) - Impala crashes randomly on different queries with GROUP BY.
- [KUDU-2238](#) - Fixed a bug where large updates were not flushed to disk, even when the server is under memory pressure..
- [KUDU-2274](#) - Fixed a very rare bug where a tombstone being copied over could produce a tablet replica in an inconsistent state.
- [KUDU-2343](#) - Java client doesn't properly reconnect to leader master when old leader is online.
- [OOZIE-3132](#) - Instrument SLACalculatorMemory
- [PARQUET-1133](#) - Add int96 support by returning bytearray, Skip originalType comparison for map types when originalType is null
- [PARQUET-1217](#) - Incorrect handling of missing values in Statistics
- [SENTRY-1459](#) - Alter view with HMS Client fails with "java.lang.IllegalArgumentException: Can not create a Path from a null string"
- [SENTRY-2124](#) - LeaderStatusMonitor.toString() throws IllegalFormatConversionException with AtomicLong
- [SOLR-11873](#) - ofand SOLR-7391: Use time based expiration cache in all necessary places in HdfsDirectoryFactory.
- [SPARK-19982](#) - Disable flaky java bean encoder tests ().
- [SPARK-23438](#) - [DSTREAMS] Fix DStreams data loss with WAL when driver crashes
- [SQOOP-2976](#) - avro padding

Issues Fixed in CDH 5.13.2

CDH 5.13.2 fixes the following issues.

Hive Replications Can Fail Intermittently

This issue, which was also reported in [TSB-2018-276](#), is fixed in CDH 5.13.2. For further details, see the original [Known Issue](#).

Secure Impala-Kudu clusters need to be restarted frequently

Due to Kudu-2264, secure, long-running Impala clusters on CDH 5.13.x and later will be unable to interact with Kudu tables after some time because the Kudu client used by Impala will no longer be able to authenticate with Kudu. The problem starts after the client's ticket can no longer be renewed, which is typically after 7 days. Restarting the Impala service will allow Impala to work with Kudu again.

Impala will begin failing to read or write Kudu tables with errors in the Impala daemon logs resembling:

```

W1121 22:47:47.231425 50416 ConnectToCluster.java:302] Error receiving response from
cdh-master-4db933ef.cdh-cluster.internal:7051
Java exception follows:
org.apache.kudu.client.NonRecoverableException: Server requires Kerberos, but this client
is not
authenticated (kinit)

at org.apache.kudu.client.Negotiator.evaluateChallenge(Negotiator.java:705)
at org.apache.kudu.client.Negotiator.sendSaslInitiate(Negotiator.java:581)
at org.apache.kudu.client.Negotiator.startAuthentication(Negotiator.java:545)
at org.apache.kudu.client.Negotiator.handleTlsMessage(Negotiator.java:499)
at org.apache.kudu.client.Negotiator.handleResponse(Negotiator.java:264)
at org.apache.kudu.client.Negotiator.messageReceived(Negotiator.java:231)
at ...
Caused by: javax.security.sasl.SaslException: GSS initiate failed [Caused by GSSException:
No valid credentials provided (Mechanism level: Failed to find any Kerberos tgt)]

at com.sun.security.sasl.gsskerb.GssKrb5Client.evaluateChallenge(GssKrb5Client.java:211)
at org.apache.kudu.client.Negotiator$1.run(Negotiator.java:691)
at org.apache.kudu.client.Negotiator$1.run(Negotiator.java:688)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:422)
at org.apache.kudu.client.Negotiator.evaluateChallenge(Negotiator.java:687)
... 35 more

Caused by: GSSException: No valid credentials provided (Mechanism level: Failed to find
any
Kerberos tgt)

at sun.security.jgss.krb5.Krb5InitCredential.getInstance(Krb5InitCredential.java:147)
at sun.security.jgss.krb5.Krb5MechFactory.getCredentialElement(Krb5MechFactory.java:122)
at sun.security.jgss.krb5.Krb5MechFactory.getMechanismContext(Krb5MechFactory.java:187)
at sun.security.jgss.GSSManagerImpl.getMechanismContext(GSSManagerImpl.java:224)
at sun.security.jgss.GSSContextImpl.initSecContext(GSSContextImpl.java:212)
at sun.security.jgss.GSSContextImpl.initSecContext(GSSContextImpl.java:179)
at com.sun.security.sasl.gsskerb.GssKrb5Client.evaluateChallenge(GssKrb5Client.java:192)
... 40 more

```

Alternatively, errors may be seen like the following:

```

I0102 11:58:36.688992 3587 jni-util.cc:196]
org.apache.impala.common.ImpalaRuntimeException: Unable to initialize the Kudu scan node

at org.apache.impala.planner.KuduScanNode.init(KuduScanNode.java:133)
at org.apache.impala.planner.SingleNodePlanner.createScanNode(SingleNodePlanner.java:1312)
at
org.apache.impala.planner.SingleNodePlanner.createTableRefNode(SingleNodePlanner.java:1519)
at
org.apache.impala.planner.SingleNodePlanner.createTableRefsPlan(SingleNodePlanner.java:779)
at
org.apache.impala.planner.SingleNodePlanner.createSelectPlan(SingleNodePlanner.java:617)
at org.apache.impala.planner.SingleNodePlanner.createQueryPlan(SingleNodePlanner.java:260)
at
org.apache.impala.planner.SingleNodePlanner.createSingleNodePlan(SingleNodePlanner.java:150)
at org.apache.impala.planner.Planner.createPlan(Planner.java:98)
at org.apache.impala.service.Frontend.createExecRequest(Frontend.java:1004)
at org.apache.impala.service.Frontend.createExecRequest(Frontend.java:1100)
at org.apache.impala.service.JniFrontend.createExecRequest(JniFrontend.java:156)

Caused by: org.apache.kudu.client.NonRecoverableException: cannot re-acquire
authentication
token after 5 attempts

at
org.apache.kudu.client.AuthnTokenReacquirer$1NewAuthnTokenErrB.failQueuedRpcs(AuthnTokenReacquirer.java:167)
at
org.apache.kudu.client.AuthnTokenReacquirer$1NewAuthnTokenErrB.call(AuthnTokenReacquirer.java:159)
at
org.apache.kudu.client.AuthnTokenReacquirer$1NewAuthnTokenErrB.call(AuthnTokenReacquirer.java:142)
at com.stumbleupon.async.Deferred.doCall(Deferred.java:1262)

```



```

at com.stumbleupon.async.Deferred.runCallbacks(Deferred.java:1241)
at com.stumbleupon.async.Deferred.callback(Deferred.java:989)
at
org.apache.kudu.client.ConnectToCluster.incrementCountAndCheckExhausted(ConnectToCluster.java:223)
at org.apache.kudu.client.ConnectToCluster.access$000(ConnectToCluster.java:48)
at
org.apache.kudu.client.ConnectToCluster$ConnectToMasterErrCB.call(ConnectToCluster.java:304)
at
org.apache.kudu.client.ConnectToCluster$ConnectToMasterErrCB.call(ConnectToCluster.java:293)
at com.stumbleupon.async.Deferred.doCall(Deferred.java:1262)
at com.stumbleupon.async.Deferred.runCallbacks(Deferred.java:1241)
at com.stumbleupon.async.Deferred.handleContinuation(Deferred.java:1297)

```

Products Affected: Impala

Releases Affected: CDH 5.13.0, CDH 5.13.1, and CDH 5.14.0

Users Affected: Users running a secure cluster with Impala and Kudu.

Severity: Medium

Impact: Impala will not be able to query or write to Kudu tables until the service is restarted. The Impala service will need to be restarted periodically in order to continue being able to work with Kudu tables.

Immediate action required: To work around the issue, users can restart the Impala service at least once per ticket renewal period. This period is typically 7 days.

Addressed in release: CDH 5.13.2, CDH 5.13.3, and CDH 5.14.2

Apache Hadoop MapReduce Job History Server (JHS) vulnerability

A vulnerability in Hadoop's Job History Server allows a cluster user to expose private files owned by the user running the MapReduce Job History Server (JHS) process. See <http://seclists.org/oss-sec/2018/q1/79> for reference.

Products affected: Apache Hadoop MapReduce

Releases affected: All releases prior to CDH 5.12.0. CDH 5.12.0, CDH 5.12.1, CDH 5.12.2, CDH 5.13.0, CDH 5.13.1, CDH 5.14.0

Users affected: Users running the MapReduce Job History Server (JHS) daemon

Date/time of detection: November 8, 2017

Detected by: Man Yue Mo of lgtm.com

Severity (Low/Medium/High): High

Impact: The vulnerability allows a cluster user to expose private files owned by the user running the MapReduce Job History Server (JHS) process. The malicious user can construct a configuration file containing XML directives that reference sensitive files on the MapReduce Job History Server (JHS) host.

CVE: CVE-2017-15713

Immediate action required: Upgrade to a release where the issue is fixed.

Addressed in release/refresh/patch: CDH 5.13.2, 5.14.2, 5.15.0

HUE

Group sync fails when a non-standard (not uid) user_name_attr is used

Group sync fails when a non-standard (not uid) user_name_attr is used. For example, if a customer user "LDAP User Attribute" is used sync fails for users where the users uid value does not match the user_name_attr defined in the configuration. There is also a "... is not JSON serializable" error when running "Sync LDAP Users/Groups"

Bug: CDH-61984

Apache Oozie Server Vulnerability

A vulnerability in the Oozie Server allows a cluster user to read private files owned by the user running the Oozie Server process.

Products affected: Oozie

Releases affected: All releases prior to CDH 5.12.0, CDH 5.12.0, CDH 5.12.1, CDH 5.12.2, CDH 5.13.0, CDH 5.13.1, CDH 5.14.0

Users affected: Users running the Oozie Server

Date/time of detection: November 13, 2017

Detected by: Daryn Sharp and Jason Lowe of Oath (formerly Yahoo! Inc)

Severity: (Low/Medium/High) High

Impact: The vulnerability allows a cluster user to read private files owned by the user running the Oozie Server process. The malicious user can construct a workflow XML file containing XML directives and configuration that reference sensitive files on the Oozie server host.

CVE: CVE-2017-15712

Immediate action required: Upgrade to release where the issue is fixed.

Addressed in release/refresh/patch: CDH 5.13.2, 5.14.2

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.13.2:

- [FLUME-3046](#) - Kafka Sink and Source Configuration Improvements
- [FLUME-3100](#) - Support arbitrary header substitution for topic of Kafka Sink
- [FLUME-3158](#) - Upgrade surefire version and config
- [FLUME-3195](#) - Split the KafkaChannelTest to avoid timeouts
- [FLUME-3201](#) - Fix SyslogUtil to handle RFC3164 format in December correctly
- [HADOOP-14687](#) - AuthenticatedURL will reuse bad/expired session cookies
- [HADOOP-14814](#) - Fix incompatible API change on FsServerDefaults to HADOOP-14104
- [HADOOP-15143](#) - NPE due to Invalid KerberosTicket in UGI
- [HDFS-11799](#) - Introduce a config to allow setting up write pipeline with fewer nodes than replication factor
- [HDFS-12052](#) - Set SWEBHDFS delegation token kind when ssl is enabled in HttpFS
- [HDFS-12485](#) - expunge may fail to remove trash from encryption zone
- [HDFS-12985](#) - NameNode crashes during restart after an OpenForWrite file present in the Snapshot got deleted.
- [MAPREDUCE-5124](#) - AM lacks flow control for task events
- [MAPREDUCE-7015](#) - Possible race condition in JHS if the job is not loaded
- [MAPREDUCE-7020](#) - Task timeout in uber mode can crash AM
- [MAPREDUCE-7028](#) - Concurrent task progress updates causing NPE in Application Master
- [YARN-4813](#) - TestRMWebServicesDelegationTokenAuthentication.testDoAs fails intermittently
- [YARN-7363](#) - ContainerLocalizer don't have a valid log4j config in case of Linux container executor.
- [YARN-7590](#) - Improve container-executor validation check
- [YARN-7622](#) - Allow fair-scheduler configuration on HDFS
- [YARN-7705](#) - Create the container log directory with correct sticky bit in C code
- [YARN-7758](#) - Add an additional check to the validity of container and application ids passed to container-executor
- [HBASE-14252](#) - RegionServers fail to start when setting hbase.ipc.server.callqueue.scan.ratio to 0
- [HBASE-15354](#) - Use same criteria for clearing meta cache for all operations
- [HBASE-16980](#) - TestRowProcessorEndpoint failing consistently
- [HBASE-18625](#) - Splitting of region with replica, doesn't update region list in serverHolding. A server crash leads to overlap.
- [HBASE-19215](#) - Incorrect exception handling on the client causes incorrect call timeouts and byte buffer allocations on the server
- [HIVE-9447](#) - Metastore: inefficient Oracle query for removing unused column descriptors when adding or dropping a table or partition
- [HIVE-11569](#) - Use PreOrderOnceWalker where feasible

- [HIVE-11609](#) - Capability to add a filter to hbase scan via composite key doesn't work
- [HIVE-14303](#) - CommonJoinOperator.checkAndGenObject should return directly to avoid NPE if ExecReducer.close is called twice
- [HIVE-15519](#) - BitSet not computed properly for ColumnBuffer subset
- [HIVE-15859](#) - HoS: Write RPC messages in event loop
- [HIVE-16708](#) - Exception while renewing a Delegation Token
- [HIVE-17022](#) - Add mode in lock debug statements
- [HIVE-17624](#) - MapredLocalTask running in separate JVM could throw ClassNotFoundException
- [HIVE-17969](#) - Metastore to alter table in batches of partitions when renaming table
- [HIVE-18009](#) - Multiple lateral view query is slow on Hive on Spark
- [HIVE-18127](#) - Do not strip '--' comments from shell commands issued from CliDriver
- [HUE-7353](#) - [search] Timeline chart shouldn't degrade the page performance
- [HUE-7564](#) - Remove one of the get_connection(ldap_config) function in ldap_access.py.
- [HUE-7567](#) - [all] Fix test desktop.lib.rest.test_clear_cookies
- [HUE-7583](#) - [editor] Export result to table throws already exists error
- [HUE-7592](#) - [jib] Hide older browser if the new ones is enabled
- [HUE-7622](#) - [jib] Reruning a wokflow redirects to a json result
- [HUE-7624](#) - [core] Support multi-authentication with AllowFirstUserDjangoBackend and LdapBackend
- [HUE-7646](#) - [sentry] Changing role privileges removes HDFS ACLs
- [HUE-7647](#) - [core] Group sync fails when a non-standard (not uid) user_name_attr is used
- [HUE-7666](#) - [editor] Presentation mode with variables shouldn't break the editor
- [HUE-7667](#) - [core] Fix http client pool by correctly mounting Transport Adapter to a prefix.
- [HUE-7676](#) - [core] skip_trash for hue_config_validation
- [HUE-7692](#) - [editor] Upgrade failure from from CDH <= 5.4
- [HUE-7697](#) - [core] Allow Kerberos Renewer to retry kinit multiple times
- [HUE-7739](#) - [metastore] Fix column context popover in the table browser
- [HUE-7755](#) - [oozie] Adding Distcp arguments and properties
- [HUE-7790](#) - [assist] Silence all API errors in the right assistant
- [HUE-7794](#) - [fb] Modify/upload file failed if S3 bucket encryption is SSE-KMS
- [HUE-7802](#) - [dashboard] Protect against missing function in old dashboards
- [HUE-7806](#) - [core] Show all the users "Switch to Hue 4" option
- [HUE-7822](#) - [core] Unable to clone oozie workflow in Hue 4.0
- [HUE-7830](#) - [core] Non supported feature use_default_configuration doesn't work in Hue 4
- [HUE-7874](#) - [oozie] SSH action missing configuration for retry-max and retry-interval
- [HUE-7909](#) - [editor] Hide options in session settings
- [HUE-7916](#) - [frontend] Avoid breaking of scatterplot chart for some grouping values
- [IMPALA-2810](#) - Remove column stats restoration when altering table.
- [IMPALA-4664](#) - Unexpected string conversion in Shell.
- [IMPALA-5167](#) - Revert "Reduce the number of Kudu clients created (FE)".
- [IMPALA-5664](#) - Unix time to timestamp conversions may crash Impala
- [IMPALA-5926](#) - Avoid printing expensive stack when closing a session.
- [IMPALA-5936](#) - Operator '%' overflows on large decimals.
- [IMPALA-5940](#) - Avoid log spew by using Status::Expected.
- [IMPALA-5951](#) - Remove flaky test_catalogd_timeout.
- [IMPALA-5987](#) - LZ4 Codec silently produces bogus compressed data for large inputs.
- [IMPALA-6004](#) - Fix the test_row_filters failure on ASAN.
- [IMPALA-6053](#) - Fix the exception when storadelds don't match hosts.
- [IMPALA-6060](#) - Check the return value of JNI exception handling functions.
- [IMPALA-6067](#) - Enable S3 access via IAM roles for EC2 VMs.
- [IMPALA-6069](#) - Fix CodegenAnyVal's handling of 'nan'.

- [IMPALA-6100](#) - Increase test_exchange_delays timeout on slow builds.
- [IMPALA-6106](#) - Handle comments before set in test parser.
- [IMPALA-6114](#) - Require type equality for NumericLiteral::localEquals().
- [IMPALA-6137](#) - Fix text scanner split delim mem mgmt.
- [IMPALA-6183](#) - Fix Decimal to Double conversion.
- [IMPALA-6184](#) - Clean up after ScalarExprEvaluator::Clone() fails.
- [IMPALA-6280](#) - Materialize TupleIsNullPredicate for insert sorts.
- [IMPALA-6284](#) - Mark the intermediate decimal avg struct as packed.
- [IMPALA-6286](#) - Remove invalid runtime filter targets.
- [IMPALA-6291](#) - Disable AVX512 codegen in LLVM.
- [IMPALA-6295](#) - Fix mix/max handling of 'nan' and 'inf'.
- [IMPALA-6307](#) - CTAS statement fails with duplicate column exception.
- [IMPALA-6308](#) - Fix bad Status() usage in data-stream-sender.cc.
- [IMPALA-6348](#) - Redact only sensitive fields in runtime profiles.
- [IMPALA-6364](#) - Bypass file handle cache for ineligible files.
- [IMPALA-6381](#) - Increase test_exchange_delays timeout for isilon.
- [IMPALA-6384](#) - RequestPoolService should honor custom group mapping config.
- [IMPALA-6388](#) - Fix the Union node number of hosts estimation.
- [IMPALA-6418](#) - Find a reliable way to detect supported TLS versions.
- [IMPALA-6435](#) - Disable codegen for CHAR literals.
- [KUDU-2147](#) - Consensus: Unknown leader should not be treated as valid UUID
- [KUDU-2231](#) - Sparse column predicate can cause excessive data-block reads.
- [KUDU-2251](#) - rowset size can overflow int in RowSetInfo.
- rpc: Allow setting --rpc_tls_min_protocol on older RHEL versions.
- [OOZIE-1401](#) - amend PurgeCommand should purge the workflow jobs w/o end_time
- [OOZIE-2045](#) - Symlink support for sharelib
- [OOZIE-3078](#) - PasswordMasker throws NPE with null arguments
- [OOZIE-3103](#) - Improve log message regarding max concurrency for coordinators
- [OOZIE-3112](#) - SparkConfigurationService overwrites properties provided via --properties-file option in SparkAction
- [PIG-5325](#) - Schema disambiguation can't be turned off for nested schemas
- [SENTRY-1712](#) - Add trigger mechanism for Sentry to push full path snapshot to Name Node.
- [SENTRY-1951](#) - Old SentryStore.retrieveFullPathsImage() should be removed
- [SENTRY-2019](#) - Improve logging in SentryPlugin
- [SENTRY-2031](#) - Add trigger mechanism for Sentry to pull full path snapshot from HMS
- [SENTRY-2032](#) - Leading Slashes need to removed when creating HMS path entries
- [SENTRY-2040](#) - When getting Snapshots from HMS we need more logging around cases when a snapshot is not being received
- [SENTRY-2046](#) - Create a full snapshot if AUTHZ_PATHS_SNAPSHOT_ID is empty, even if HMS and Sentry Notifications are in sync
- [SENTRY-2047](#) - isEmptyCore method in SentryStore has references to MAuthzPathsMapping when it should be generic
- [SENTRY-2066](#) - DB name is not set for AlterTable
- [SENTRY-2078](#) - Have sentry server print an obvious INFO level log message when it becomes the writer
- [SENTRY-2120](#) - Escape input string for error response message in LogLevelServlet
- [SPARK-16845](#) - [SQL][BRANCH-1.6] GeneratedClass\$SpecificOrdering` grows beyond 64 KB
- [SPARK-21551](#) - [PYTHON] Increase timeout for PythonRDD.serveIterator
- [SQOOP-3196](#) - Modify MySQLAuthTest to use configurable test database parameters
- [SQOOP-3197](#) - Database credentials cannot be configured in PostgreSQL third party tests
- [SQOOP-3229](#) - Document how to run third party tests manually with databases running in docker
- [SQOOP-3233](#) - SqoopHCatImportHelper.convertNumberTypes check for Varchar instead of Char

- [SQOOP-3241](#) - ImportAllTablesTool uses the same SqoopOptions object for every table import
- [ZOOKEEPER-2249](#) - CRC check failed when preAllocSize smaller than node data

Issues Fixed in CDH 5.13.1

CDH 5.13.1 fixes the following issues.

HDFS

Cannot re-encrypt an encryption zone if a previous re-encryption on it was canceled

When canceling a re-encryption on an encryption zone, the status of the re-encryption would continue to show "Processing". When this occurred, future re-encrypt commands for this encryption zone failed inside the NameNode, and the re-encryption never completed.

Bug: HDFS-12518, CDH-59073

Upstream Issues Fixed in 5.13.1

The following upstream issues are fixed in CDH 5.13.1:

- [FLUME-3186](#) - Make asyncHbaseClient config parameters available from Flume config
- [HADOOP-13017](#) - Implementations of InputStream.read(buffer, offset, bytes) to exit 0 if bytes==0
- [HDFS-6763](#) - Initialize file system-wide quota once on transitioning to active
- [HDFS-7582](#) - Enforce maximum number of ACL entries separately per access and default.
- [HDFS-8540](#) - Mover should exit with NO_MOVE_BLOCK if no block can be moved
- [HDFS-8541](#) - Mover should exit with NO_MOVE_PROGRESS if there is no move progress
- [HDFS-8865](#) - Improve quota initialization performance
- [HDFS-9003](#) - ForkJoin thread pool leaks
- [HDFS-9396](#) - Total files and directories on jmx and web UI on standby is uninitialized
- [HDFS-11384](#) - Balancer disperses getBlocks calls to avoid NameNode's rpc queue saturation
- [HDFS-11742](#) - Improve balancer usability after HDFS-8818
- [HDFS-11755](#) - Underconstruction blocks can be considered missing
- [HDFS-12424](#) - Datatable sorting on the Datanode Information page in the Namenode UI is broken
- [HDFS-12494](#) - libhdfs SIGSEGV in setTLSEceptionStrings
- [HDFS-12518](#) - Re-encryption should handle task cancellation and progress better.
- [HDFS-12642](#) - Log block and datanode details in BlockRecoveryWorker.
- [YARN-6757](#) - Refactor the usage of yarn.nodemanager.linux-container-executor.cgroups.mount-path.
- [YARN-7099](#) - ResourceHandlerModule.parseConfiguredCGroupPath only works for privileged yarn users.
- [YARN-7207](#) - Cache the RM proxy server address.
- [YARN-7261](#) - Add debug message for better download latency monitoring.
- [HBASE-18665](#) - ReversedScannerCallable invokes getRegionLocations incorrectly
- [HBASE-18675](#) - Making {max,min}SessionTimeout configurable for MiniZooKeeperCluster
- [HBASE-18924](#) - Backport HBASE-18568 (Correct metric of numRegions) to branch-1.2 and branch-1.3
- [HIVE-12742](#) - NULL table comparison within CASE does not work as it did in previous Hive versions.
- [HIVE-16758](#) - Better Select Number of Replications
- [HIVE-16784](#) - Missing lineage information when hive.blobstore.optimizations.enabled is true
- [HIVE-17218](#) - Canonical-ize hostnames for Hive metastore, and HS2 servers.
- [HIVE-17274](#) - RowContainer spills for timestamp column throws exception
- [HIVE-17640](#) - Comparison of date return null if time part is provided in string.
- [HIVE-17659](#) - get_token thrift call fails for DBTokenStore in remote HMS mode
- [HIVE-17805](#) - SchemaTool validate locations should not return exit 1
- [HUE-6835](#) - [autocomplete] Update the Impala autocompleter to match the latest Impala version
- [HUE-6851](#) - [metastore] Loading a table is missing loading feedback
- [HUE-7088](#) - [core] Display a warning when users access a Hue 3 page while on Hue 4
- [HUE-7124](#) - [frontend] The global search should highlight the match for Hue documents

- [HUE-7182](#) - [editor] Accented characters not inserted correctly
- [HUE-7222](#) - [desktop] Document filtering on home page always return an empty list
- [HUE-7229](#) - [editor] Expanding a result row a second time does not reset the scroll of the modal to the top
- [HUE-7230](#) - [editor] In full screen the left column filtering is too high and not shown when scrolling down
- [HUE-7237](#) - [job] Doc link in Oozie workflow has a mini scrollbar
- [HUE-7246](#) - [importer] Input file input styling can be too short sometimes
- [HUE-7249](#) - [indexer] Index sample data tab can be empty when the collection is small
- [HUE-7259](#) - [autocomplete] Suggest REGEXP instead of REGEX
- [HUE-7260](#) - [assist] Disable "Open in importer" option for folders in HDFS
- [HUE-7264](#) - [core] Fix cx_Oracle removed numbersAsStrings in 6.0
- [HUE-7270](#) - [core] Do not display warning in Hue 3 on login page
- [HUE-7276](#) - [autocomplete] Support variable references in limit clause for Hive
- [HUE-7277](#) - [autocomplete] Fix issue where columns are marked as tables in the select list when qualified table references are used
- [HUE-7279](#) - [autocomplete] Mark column alias locations
- [HUE-7283](#) - [editor] Fix issue where analytic functions aren't found in the context popover
- [HUE-7285](#) - [oozie] Coordinator last modification date is always empty
- [HUE-7288](#) - [editor] Don't check for risks for incomplete statements
- [HUE-7324](#) - [autocomplete] Prevent js error on column alias selection in the autocomplete dropdown
- [HUE-7325](#) - [assist] Gracefully handle missing tables or backend errors in the assistant panel
- [HUE-7326](#) - [assist] Include table names in the right assistant filter
- [HUE-7328](#) - [assist] Don't reload assistant tables on every change
- [HUE-7331](#) - [assist] Make the assistant panel aware of snippet database changes
- [HUE-7335](#) - [editor] Make it possible to close the editor sessions panel form within the panel
- [HUE-7390](#) - [autocomplete] Fall back to only autocomplete the active statement in case there are surrounding errors
- [HUE-7398](#) - [dashboard] Do not 500 on Dashboard opening when the first collection is not queryable
- [HUE-7402](#) - [core] Truncate SEARCH Column in Document2 for Oracle database
- [HUE-7412](#) - [editor] Fix indexes for chart limit array slice
- [HUE-7427](#) - [editor] Keep the selected database after execution of a statement that refreshes the assist panel
- [HUE-7428](#) - [core] Opening a view directly from top search lands on a 404
- [HUE-7430](#) - [core] Update translations files
- [HUE-7438](#) - [editor] Query execution time is not updated on short queries
- [HUE-7439](#) - [editor] Move optimizer call to after check status
- [HUE-7446](#) - [core] RM HA gets called by mistake on each navopt call
- [HUE-7456](#) - [core] Do not retry Thrift call on read operation timed out from ssl.SSLError
- [HUE-7457](#) - [assist] Don't fetch top columns for tables in the right assistant
- [HUE-7465](#) - [assist] Fetch right assistant columns synchronously to reduce backend load
- [HUE-7470](#) - [editor] Only check for risks when a snippet has focus
- [HUE-7476](#) - [core] Add Hue shell command to set CM environment
- [HUE-7477](#) - [frontend] Close the left nav on click
- [HUE-7489](#) - [editor] CloseOperation when GetOperationStatus is in an ERROR_STATE
- [HUE-7519](#) - [assist] Fix JS error when the sample tab can't load samples
- [HUE-7520](#) - [editor] Prevent JS error from being thrown when closing the session
- [HUE-7521](#) - [security] Prevent JS error on printAcl not defined
- [HUE-7522](#) - [frontend] Link to the latest editor from the oldest editor alert
- [HUE-7525](#) - [home] Prevent js error when dragging to select documents
- [HUE-7526](#) - [frontend] Avoid XSS in the wysiwyg5 editor and fix the binding in Hue 4
- [HUE-7527](#) - [assist] Prevent scroll on drag in the foreachVisible binding
- [HUE-7530](#) - [editor] Avoid red error popup on close of expired query

- [HUE-7531](#) - [frontend] Allow double click to see detail on old dataTables too
- [HUE-7533](#) - [fb] Wrong error message in UI with 500 error when user tries to edit read_only file in File browser
- [HUE-7534](#) - [editor] Fix IE js error on missing endsWith in the syntax checker
- [HUE-7535](#) - [editor] Fix issue where the syntax checker marks incomplete statements as complete
- [HUE-7536](#) - [frontend] Avoid XSS in the table row display modal, HueDatatable and TableExtender2
- [HUE-7537](#) - [editor] Make SqlFunctions optional for the parsers
- [HUE-7538](#) - [frontend] Override default knockout html binding to prevent XSS
- [HUE-7543](#) - [editor] Prevent the risk icon from moving the editor down
- [HUE-7552](#) - [fb] Copy Folder Into Itself Causes Endless Loop
- [HUE-7567](#) - [all] Add default timeout to REST calls
- [IMPALA-4682](#) - Remove Preconditions check from analyzeAggregation().
- [IMPALA-4951](#) - Fix database visibility for user with only column privilege
- [IMPALA-5597](#) - Try casting targetExpr when building runtime filter plan
- [IMPALA-5856](#) - Fix outer join predicate assignment.
- [IMPALA-5923](#) - Print binary ID as hex in ChildQuery::Cancel()
- [IMPALA-5954](#) - Set DO_NOT_UPDATE_STATS in alterTable() RPCs to HMS.
- [IMPALA-5955](#) - Use totalSize tblproperty instead of rawDataSize.
- [IMPALA-5983](#) - Fix crash in to/from_utc_timestamp
- [KUDU-1788](#) - Increase Raft RPC timeout to 30sec to avoid fruitless retries.
- [KUDU-2130](#) - (part 2): more fixes for ITClientStress
- [KUDU-2130](#) - java client: handle termination during negotiation edge case
- [KUDU-2167](#) - fix C++ client crash due to bad assumption regarding scan data
- [KUDU-2170](#) - Masters can start with duplicates specified
- [KUDU-2173](#) - Partitions are incorrectly pruned when range-partitioned on a PK prefix
- [KUDU-2188](#) - restore Java 7 compatibility to artifacts built with JDK8
- [KUDU-2209](#) - HybridClock doesn't handle changes in STA_NANO flag
- [SENTRY-1906](#) - Sentry clients to retry connections to server with delay to avoid failing fast
- [SENTRY-1966](#) - Improve logging of HMS sync data (paths and permissions) flowing from Sentry to NameNode
- [SENTRY-1993](#) - StringIndexOutOfBoundsException in HMSPathsDumper.java
- [SENTRY-2027](#) - backportCreate mechanism of delivering commands via WebUI
- [SPARK-17460](#) - [SQL] Make sure sizeInBytes in Statistics will not overflow
- [SPARK-17511](#) - Yarn Dynamic Allocation: Avoid marking released container as Failed
- [SPARK-17531](#) - Fix the "hive.metastore.uris" value that was changed by the backport of 1.6 branches.
- [SQOOP-3014](#) - Sqoop with HCatalog import loose precision for large numbers that does not fit into double
- [SQOOP-3243](#) - Importing BLOB data causes 'Stream closed' error on encrypted HDFS
- [ZOOKEEPER-2347](#) - Deadlock shutting down zookeeper
- [ZOOKEEPER-2726](#) - backportPatch for ZOOKEEPER-2693 introduces potential race condition

Issues Fixed in CDH 5.13.0

CDH 5.13.0 fixes the following issues.

Sqoop 1

Sqoop does not support Hive metastore HA for Parquet format

If the Hive metastore (HMS) is configured for high availability (HA), the value of the `hive.metastore.uris` property in `/etc/hive/conf.cloudera.HIVE-1/hive-site.xml` contains a list of URIs instead of one URI. For example:

```
<value>thrift://hive-1.example.com:9083,thrift://hive-2.example.com:9083,thrift://hive-3.example.com:9083</value>
```

Sqoop uses Kite for Parquet file handling and Kite does not support HMS HA. Using Kite, Sqoop takes the first URI from the `hive.metastore.uris` list and uses that to connect to the HMS. If that address is not available, the Sqoop command fails.

This issue occurs when you use the following Sqoop arguments with the `sqoop-import` command: `--hive-import` and `--as-parquetfile`.

Bug: CDH-57895

Upstream Issues Fixed in 5.13.0

The following upstream issues are fixed in CDH 5.13.0:

- [AVRO-1985](#) - DataFileWriter.appendTo(SeekableInput, OutputStream) does not close the input anymore
- [FLUME-2752](#) - Fix AvroSource startup resource leaks
- [FLUME-2905](#) - Fix NetcatSource file descriptor leak if startup fails
- [FLUME-3093](#) - Groundwork for version changes in root pom
- [FLUME-3144](#) - Improve Log4jAppender's performance by allowing logging collection of messages
- [FLUME-3152](#) - Add Flume Metric for Backup Checkpoint Errors
- [HADOOP-9137](#) - Support connection limiting in IPC server
- [HADOOP-11572](#) - s3a delete() operation fails during a concurrent delete of child entries
- [HADOOP-13414](#) - Hide Jetty Server version header in HTTP responses
- [HADOOP-13588](#) - ConfServlet should respect Accept request header
- [HADOOP-13628](#) - Support to retrieve specific property from configuration via REST API
- [HADOOP-13700](#) - Remove unthrown IOException from TrashPolicy#initialize and #getInstance signatures.
- [HADOOP-13827](#) - Add reencryptEncryptedKey interface to KMS.
- [HADOOP-13867](#) - FilterFileSystem should override rename(.., options) to take effect of Rename options called via FilterFileSystem implementations
- [HADOOP-14024](#) - KMS JMX endpoint throws ClassNotFoundException
- [HADOOP-14038](#) - Rename ADLS credential properties
- [HADOOP-14174](#) - Set default ADLS access token provider type to ClientCredential
- [HADOOP-14251](#) - Credential provider should handle property key deprecation
- [HADOOP-14260](#) - Configuration.dumpConfiguration should redact sensitive information
- [HADOOP-14324](#) - Refine S3 server-side-encryption key as encryption secret; improve error reporting and diagnostics
- [HADOOP-14349](#) - Rename ADLS CONTRACT_ENABLE_KEY
- [HADOOP-14376](#) - Memory leak when reading a compressed file using the native library
- [HADOOP-14401](#) - maven-project-info-reports-plugin can be removed
- [HADOOP-14407](#) - DistCp - Introduce a configurable copy buffer size.
- [HADOOP-14440](#) - Add metrics for connections dropped
- [HADOOP-14464](#) - hadoop-aws doc header warning #5 line wrapped
- [HADOOP-14504](#) - ProvidedFileStatusIterator#next() may throw IndexOutOfBoundsException
- [HADOOP-14511](#) - WritableRpcEngine.Invocation#toString NPE on null parameters
- [HADOOP-14515](#) - Addendum. Specifically configure zookeeper-related log levels in KMS log4j.
- [HADOOP-14515](#) - Specifically configure zookeeper-related log levels in KMS log4j.
- [HADOOP-14521](#) - Revert "KMS client needs retry logic"
- [HADOOP-14521](#) - KMS client needs retry logic
- [HADOOP-14523](#) - OpensslAesCtrCryptoCodec.finalize() holds excessive amounts of memory
- [HADOOP-14524](#) - Make CryptoCodec Closeable so it can be cleaned up proactively.
- [HADOOP-14533](#) - Size of args cannot be less than zero in TraceAdmin#run as its linkedlist
- [HADOOP-14540](#) - Replace MRV1 specific terms in HostsFileReader
- [HADOOP-14542](#) - Add IOUtils.cleanupWithLogger that accepts slf4j logger API
- [HADOOP-14563](#) - LoadBalancingKMSClientProvider#warmUpEncryptedKeys swallows IOException
- [HADOOP-14586](#) - StringIndexOutOfBoundsException breaks org.apache.hadoop.util.Shell on 2.7.x with Java 9
- [HADOOP-14627](#) - Support MSI and DeviceCode token provider in ADLS
- [HADOOP-14678](#) - AdlFileSystem#initialize swallows exception when getting user name
- [HADOOP-14688](#) - Intern strings in KeyVersion and EncryptedKeyVersion
- [HADOOP-14705](#) - Add batched interface reencryptEncryptedKeys to KMS.

- [HADOOP-14760](#) - Add missing override to LoadBalancingKMSSClientProvider.
- [HADOOP-14824](#) - Update ADLS SDK to 2.2.2 for MSI fix
- [HADOOP-14844](#) - Remove requirement to specify TenantGuid for MSI Token Provider
- [HDFS-7847](#) - Modify NNThroughputBenchmark to be able to operate on a remote NameNode
- [HDFS-8312](#) - Added permission check for moving file to Trash.
- [HDFS-8586](#) - Dead Datanode is allocated for write when client is from deadnode
- [HDFS-8674](#) - Improve performance of postponed block scans
- [HDFS-8818](#) - Changes the global moveExecutor to per datanode executors and changes MAX_SIZE_TO_MOVE to be configurable.
- [HDFS-8856](#) - Make LeaseManager#countPath O(1).
- [HDFS-8870](#) - Lease is leaked on write failure
- [HDFS-9503](#) - Use generic option -fs for NNThroughputBenchmark instead of -namenode
- [HDFS-10468](#) - HDFS read ends up ignoring an interrupt
- [HDFS-10480](#) - Add an admin command to list currently open files
- [HDFS-10506](#) - OIV's ReverseXML processor cannot reconstruct some snapshot details
- [HDFS-10899](#) - Add functionality to re-encrypt EDEKs.
- [HDFS-11210](#) - Enhance key rolling to guarantee new KeyVersion is returned from generateEncryptedKeys after a key is rolled.
- [HDFS-11303](#) - Hedged read might hang infinitely if read data from all DN failed
- [HDFS-11345](#) - Document the configuration key for FSNamesystem lock fairness
- [HDFS-11352](#) - Potential deadlock in NN when failing over
- [HDFS-11377](#) - Balancer hung due to no available mover threads
- [HDFS-11472](#) - Fix inconsistent replica size after a data pipeline failure
- [HDFS-11674](#) - reserveSpaceForReplicas is not released if append request failed due to mirror down and replica recovered
- [HDFS-11708](#) - Positional read will fail if replicas moved to different DNs after stream is opened
- [HDFS-11711](#) - DN should not delete the block On "Too many open files" Exception
- [HDFS-11738](#) - Hedged pread takes more time when block moved from initial locations
- [HDFS-11741](#) - Long running balancer may fail due to expired DataEncryptionKey
- [HDFS-11833](#) - HDFS architecture documentation describes outdated placement policy
- [HDFS-11849](#) - JournalNode startup failure exception should be logged in log file
- [HDFS-11851](#) - getGlobalJNIEEnv() may deadlock if exception is thrown
- [HDFS-11861](#) - ipc.Client.Connection#sendRpcRequest should log request name
- [HDFS-11881](#) - NameNode consumes a lot of memory for snapshot diff report generation
- [HDFS-11891](#) - DU#refresh should print the path of the directory when an exception is caught
- [HDFS-11932](#) - BPSERVICEACTOR thread name is not correctly set
- [HDFS-11945](#) - Internal lease recovery may not be retried for a long time
- [HDFS-11960](#) - Successfully closed files can stay under-replicated
- [HDFS-12042](#) - Lazy initialize AbstractINodeDiffList#diffs for snapshots to reduce memory consumption
- [HDFS-12078](#) - Add time unit to the description of property dfs.namenode.stale.datanode.interval in hdfs-default.xml
- [HDFS-12089](#) - Fix ambiguous NN retry log message in WebHDFS
- [HDFS-12117](#) - HttpFS does not seem to support SNAPSHOT related methods for WebHDFS REST Interface
- [HDFS-12139](#) - HTTPFS liststatus returns incorrect pathSuffix for path of file
- [HDFS-12140](#) - Remove BPOfferService lock contention to get block pool id
- [HDFS-12157](#) - Do fsyncDirectory(..) outside of FSDataset lock
- [HDFS-12191](#) - Provide option to not capture the accessTime change of a file to snapshot if no other modification has been done to this file
- [HDFS-12217](#) - HDFS snapshots doesn't capture all open files when one of the open files is deleted.
- [HDFS-12278](#) - LeaseManager operations are inefficient in 2.8
- [HDFS-12336](#) - Listing encryption zones still fails when deleted EZ is not a direct child of snapshottable directory

- [HDFS-12357](#) - Let NameNode to bypass external attribute provider for configured users
- [HDFS-12359](#) - Re-encryption should operate with minimum KMS ACL requirements.
- [HDFS-12363](#) - Possible NPE in BlockManager\$StorageInfoDefragmenter#scanAndCompactStorages
- [HDFS-12369](#) - Edit log corruption due to hard lease recovery of not-closed file which has snapshots.
- [HDFS-12383](#) - Re-encryption updater should handle canceled tasks better.
- [HDFS-12400](#) - Provide a way for NN to drain the local key cache before re-encryption.
- [YARN-2780](#) - Log aggregated resource allocation in rm-appsummary.log
- [YARN-3006](#) - Improve the error message when attempting manual failover with auto-failover enabled.
- [YARN-4995](#) - FairScheduler: Display per-queue demand on the scheduler page.
- [YARN-6432](#) - Revert "FairScheduler: Reserve preempted resources for corresponding applications.
- [YARN-6751](#) - Display reserved resources in web UI per queue
- [YARN-6752](#) - Display reserved resources in web UI per application
- [YARN-6802](#) - Add max AM resource and AM resource usage to leaf queue view in FairScheduler WebUI.
- [YARN-7099](#) - ResourceHandlerModule.parseConfiguredCGroupPath only works for privileged yarn users.
- [HBASE-9393](#) - Hbase does not close a dead connection with the datanode.
- [HBASE-12770](#) - Don't transfer all the queued hlogs of a dead server to the same alive server
- [HBASE-15548](#) - SyncTable: sourceHashDir is supposed to be optional but doesn't work without it
- [HBASE-15691](#) - ConcurrentModificationException in BucketAllocator
- [HBASE-15720](#) - Print row locks at the debug dump page
- [HBASE-16033](#) - Add more details in logging of responseTooSlow/TooLarge
- [HBASE-16045](#) - endtime argument for VerifyReplication was incorrectly specified in usage
- [HBASE-16090](#) - ResultScanner is not closed in SyncTable#finishRemainingHashRanges()
- [HBASE-16351](#) - Improve error reporting for license failures
- [HBASE-16409](#) - Row key for bad row should be properly delimited in VerifyReplication
- [HBASE-16423](#) - Add re-compare option to VerifyReplication to avoid occasional inconsistent rows
- [HBASE-16638](#) - Reduce the number of Connection's created in classes of hbase-spark module
- [HBASE-16772](#) - Add verbose option to VerifyReplication for logging good rows
- [HBASE-17047](#) - Add an API to get HBase connection cache statistics
- [HBASE-17131](#) - Avoid livelock caused by HRegion#processRowsWithLocks
- [HBASE-17287](#) - Master becomes a zombie if filesystem object closes
- [HBASE-17341](#) - Add a timeout during replication endpoint termination
- [HBASE-17352](#) - Fix hbase-assembly build with bash 4
- [HBASE-17534](#) - Avoid re-wrapping IOExceptions as IOExceptions
- [HBASE-17587](#) - Do not Rethrow DoNotRetryIOException as UnknownScannerException
- [HBASE-17648](#) - HBase Table-level synchronization fails between two secured(kerberized) cluster
- [HBASE-17698](#) - ReplicationEndpoint choosing sinks
- [HBASE-17773](#) - VerifyReplication tool wrongly emits Invalid arguments error
- [HBASE-17904](#) - GET runs into NoSuchElementException when using Read Replica, with hbase.
ipc.client.specificThreadForWriting to be true and hbase.rpc.client.impl to be
org.apache.hadoop.hbase.ipc.RpcClientImpl
- [HBASE-18005](#) - read replica: handle the case that region server hosting both primary replica and meta region is
down
- [HBASE-18021](#) - Add more info in timed out RetriesExhaustedException for read replica client get processing
- [HBASE-18024](#) - HRegion#initializeRegionInternals should not re-create .hregioninfo file when the region directory
no longer exists
- [HBASE-18025](#) - CatalogJanitor should collect outdated RegionStates from the AM
- [HBASE-18030](#) - Per Cell TTL tags may get duplicated with increments/Append causing tags length overflow.
- [HBASE-18035](#) - Meta replica does not give any primaryOperationTimeout to primary meta region
- [HBASE-18036](#) - Data locality is not maintained after cluster restart or SSH
- [HBASE-18066](#) - Get with closest_row_before on hbase:meta can return empty Cell during region merge/split

- [HBASE-18077](#) - Update JUnit licensing to use EPL
- [HBASE-18081](#) - The way we process connection preamble in SimpleRpcServer is broken
- [HBASE-18122](#) - Scanner id should include ServerName of region server
- [HBASE-18137](#) - Replication gets stuck for empty WALs
- [HBASE-18192](#) - Replication drops recovered queues on region server shutdown
- [HBASE-18212](#) - reduce log level for unbuffer warning.
- [HBASE-18230](#) - Remove Velocity vars from LICENSE
- [HBASE-18247](#) - Hbck to fix the case that replica region shows as key in the meta table
- [HBASE-18255](#) - Set JVM code cache size in provided configuration
- [HBASE-18358](#) - Backport HBASE-18099 'FlushSnapshotSubprocedure should wait for concurrent Region#flush() to finish'
- [HBASE-18362](#) - hbck should not report split replica parent region from meta as errors
- [HBASE-18390](#) - Sleep too long when finding region location failed
- [HBASE-18479](#) - should apply HBASE-18255 to HBASE_MASTER_OPTS too
- [HBASE-18487](#) - Minor fixes in row lock implementation
- [HIVE-6727](#) - Table level stats for external tables are set incorrectly
- [HIVE-8838](#) - Support Parquet through HCatalog
- [HIVE-9152](#) - Dynamic Partition Pruning support added
- [HIVE-9567](#) - JSON SerDe not escaping special chars when writing char/varchar data
- [HIVE-10209](#) - FetchTask with VC may fail because ExecMapper.done is true
- [HIVE-10616](#) - TypeInfoUtils doesn't handle DECIMAL with just precision specified
- [HIVE-11193](#) - ConstantPropagateProcCtx should use a Set instead of a List to hold operators to be deleted
- [HIVE-11240](#) - Change value type from int to long for HiveConf.ConfVars.METASTORESERVERMAXMESSAGEIZE
- [HIVE-11462](#) - GenericUDFStruct should constant fold at compile time
- [HIVE-11592](#) - ORC metadata section can sometimes exceed protobuf message size limit
- [HIVE-12270](#) - Add DBTokenStore support to HS2 delegation token
- [HIVE-12274](#) - Increase width of columns used for general configuration in the metastore
- [HIVE-12551](#) - Fix several kryo exceptions in branch-1
- [HIVE-12667](#) - DPP: UDFs on the partition column side does not evaluate correctly
- [HIVE-12762](#) - Common join on parquet tables returns incorrect result when hive.optimize.index.filter set to true
- [HIVE-13330](#) - ORC vectorized string dictionary reader does not differentiate null vs empty string dictionary
- [HIVE-13588](#) - NPE is thrown from MapredLocalTask.executeInChildVM.
- [HIVE-13947](#) - HoS prints wrong number for hash table size in map join scenario
- [HIVE-14178](#) - Hive::needsToCopy should reuse FileUtils::equalsFileSystem
- [HIVE-15068](#) - Run ClearDanglingScratchDir periodically inside HS2
- [HIVE-15122](#) - Hive: Upcasting types should not obscure stats
- [HIVE-15208](#) - Query string should be HTML encoded for Web UI
- [HIVE-15233](#) - UDF UUID() should be non-deterministic
- [HIVE-15428](#) - HoS DPP doesn't remove cyclic dependency
- [HIVE-15792](#) - Hive should raise SemanticException when LPAD/RPAD pad character's length is 0
- [HIVE-16183](#) - Fix potential thread safety issues with static variables
- [HIVE-16291](#) - Hive fails when unions a parquet table with itself
- [HIVE-16450](#) - Some metastore operations are not retried even with desired underlining exceptions
- [HIVE-16559](#) - Parquet schema evolution for partitioned tables may break if table and partition serdes differ
- [HIVE-16613](#) - SaslClientHandler.sendHello fails with SASL timeout instead of real cause
- [HIVE-16675](#) - Fix ConcurrentModificationException in SparkClientImpl#startDriver
- [HIVE-16729](#) - Improve location validator to check for blank paths
- [HIVE-16845](#) - INSERT OVERWRITE a table with dynamic partitions on S3 fails with NPE
- [HIVE-16869](#) - Hive returns wrong result when predicates on non-existing columns are pushed down to Parquet reader

- [HIVE-16875](#) - Query against view with partitioned child on HoS fails with privilege exception.
- [HIVE-16930](#) - HoS should verify the value of Kerberos principal and keytab file before adding them to spark-submit command parameters
- [HIVE-16935](#) - Hive should strip comments from input before choosing which CommandProcessor to run.
- [HIVE-16949](#) - Leak of threads from Get-Input-Paths and Get-Input-Summary thread pool
- [HIVE-16960](#) - Hive throws an ugly error exception when HDFS sticky bit is set
- [HIVE-16961](#) - Hive on Spark leaks spark application in case user cancels query and closes session
- [HIVE-16974](#) - Change the sort key for the schema tool validator to be ID column
- [HIVE-16978](#) - HoS: add current thread ID to the log redirector for the RemoteDriver
- [HIVE-16982](#) - WebUI "Show Query" tab prints "UNKNOWN" instead of explaining configuration option
- [HIVE-16998](#) - Add config to enable HoS DPP only for map-joins
- [HIVE-17008](#) - Fix boolean flag switchup in DropTableEvent
- [HIVE-17050](#) - Multiline queries that have comment in middle fail when executed via "beeline -e"
- [HIVE-17052](#) - Remove logging of predicate filters
- [HIVE-17087](#) - Remove unnecessary HoS DPP trees during map-join conversion
- [HIVE-17131](#) - Add InterfaceAudience and InterfaceStability annotations for SerDe APIs
- [HIVE-17132](#) - Add InterfaceAudience and InterfaceStability annotations for UDF APIs
- [HIVE-17149](#) - Hdfs directory is not cleared if partition creation failed on HMS
- [HIVE-17150](#) - CREATE INDEX execute HMS out-of-transaction listener calls inside a transaction
- [HIVE-17191](#) - Add InterfaceAudience and InterfaceStability annotations for StorageHandler APIs
- [HIVE-17225](#) - HoS DPP pruning sink ops can target parallel work objects
- [HIVE-17237](#) - HMS wastes 26.4% of memory due to dup strings in metastore.api.Partition.parameters
- [HIVE-17247](#) - HoS DPP: UDFs on the partition column side does not evaluate correctly
- [HIVE-17268](#) - WebUI / QueryPlan: query plan is sometimes null when explain output conf is on
- [HIVE-17405](#) - HoS DPP ConstantPropagate should use ConstantPropagateOption.SHORTCUT
- [HUE-1199](#) - [metastore] Wrong time stamp for all export query result to directory or new table
- [HUE-2961](#) - [indexer] Turn on the new interface by default
- [HUE-3228](#) - [editor] SQL Dashboard enhancements
- [HUE-3797](#) - [editor] Open submitted schedule in mini job browser
- [HUE-4031](#) - [editor] Smart syntax error highlighting
- [HUE-4907](#) - [editor] Provide a default value to parameters
- [HUE-5116](#) - [indexes] Remove unfinished create collection from a file
- [HUE-5217](#) - [backend] Resolve LGPL copyleft issue - Paramiko in boto 2.46.1
- [HUE-5304](#) - [search] List all the Solr cloud collections
- [HUE-5396](#) - [spark] Hide extra add property dropdown
- [HUE-5463](#) - [autocomplete] Add autocomplete support for partitions in Hive describe table statements
- [HUE-5504](#) - [oozie] Only use JDBC URL from hive2 action when hardcoded
- [HUE-5736](#) - [autocomplete] Cancel any running API requests when closing the autocompleter
- [HUE-5991](#) - [metastore] New table with Date time works with Hive but not with Impala
- [HUE-6045](#) - [Pig] Update LOG_END_PATTERN to identify completed jobs
- [HUE-6063](#) - [frontend] Remove double class attribute from element
- [HUE-6149](#) - [search] Smarter catching of Solr error message
- [HUE-6256](#) - [presentation] Result grid not displaying the cells when scrolling
- [HUE-6347](#) - [presentation] Convert player mode to presentation
- [HUE-6364](#) - [importer] Dock the bottom bar to the bottom
- [HUE-6370](#) - Revert "[jobsub] Save jobsub doesn't redirect to list-designs page."
- [HUE-6377](#) - [importer] Submit doesn't give a feedback if the call takes a long time
- [HUE-6380](#) - [editor] Move clipboard option above Save menu
- [HUE-6486](#) - Move Hue Krb ticket cache out of /tmp
- [HUE-6575](#) - [home] Convert all links to use the hueLink binding

- [HUE-6595](#) - [fb] Opening a job might throw a 500 error if its tasks are None
- [HUE-6659](#) - [import] Support RDMBs as input via Sqoop
- [HUE-6684](#) - [editor] Horizontal scrollbar of the ace editor isn't updated when the right assist is resized
- [HUE-6701](#) - [editor] The right panel resize bar gets hidden in Hue 3 on window resize
- [HUE-6702](#) - [about] About page is accessible without authentication
- [HUE-6704](#) - [fb] Handle empty page error on filebrowser filter requests
- [HUE-6705](#) - [fb] Opening directory after opening editor does not load at first click
- [HUE-6707](#) - [search] Do not send search call when the collection name is null
- [HUE-6709](#) - [editor] Column list can jump back and forth depending on its size
- [HUE-6710](#) - [notebook] Application reachable directly by users without granted access
- [HUE-6711](#) - [frontend] Render 403 page properly on Hue 4
- [HUE-6714](#) - [job] No tasks found for job when refreshing MapReduce job page
- [HUE-6719](#) - [assist] Add a close button to the assist function panel
- [HUE-6722](#) - [job] Task progress shows as 1 and not 100%
- [HUE-6723](#) - [job] Provide duration time of running Yarn jobs
- [HUE-6724](#) - [job] Job progress can have too many decimals
- [HUE-6727](#) - [job] 0/1 map progress percent gets listed as 100%
- [HUE-6728](#) - [assist] Fix truncated input in the assist filter
- [HUE-6730](#) - [fb] Use SkipFile instead of StopUpload in case of exception during file upload
- [HUE-6732](#) - [job] Add clearable to job filters
- [HUE-6735](#) - [job] A suspended workflow should also be killable
- [HUE-6737](#) - [fb] Submit workflow.xml sends to Hue 3
- [HUE-6738](#) - [job] Kill button of workflow enabled when clicking on re-run
- [HUE-6740](#) - [assist] Enable text filter for document type in the assist
- [HUE-6742](#) - [editor] Incorrect Sorting With TIMESTAMP Column Type
- [HUE-6745](#) - [frontend] The title of the browser is Filebrowser even if I'm in editor
- [HUE-6746](#) - [autocomplete] Add USER and TIMESTAMP to non-reserved keywords list
- [HUE-6748](#) - [editor] Select and copy the query from the query history
- [HUE-6749](#) - [editor] From metastore to editor with past query ran and scrolled can garble grid
- [HUE-6754](#) - [assist] Add open in dashboard to collections
- [HUE-6755](#) - [assist] Add link to create a collection
- [HUE-6756](#) - [home] The text filter should honor the active type
- [HUE-6757](#) - [fb] Used to have a fixed action menu when scrolling down
- [HUE-6758](#) - [job] Open a workflow workspace breaks the history
- [HUE-6759](#) - [job] Variable section of workflow is shown even if empty
- [HUE-6760](#) - [job] Filtering running jobs and batch killing them do not update their status
- [HUE-6762](#) - [assist] Offer right click to open table in editor
- [HUE-6763](#) - [search] Add global default if in cloud mode
- [HUE-6767](#) - [job] Submit a workflow sometimes generates a JS error
- [HUE-6768](#) - [frontend] Remove shown tooltips on app change in Hue 4
- [HUE-6770](#) - [metastore] Remove unused partitions mako page
- [HUE-6771](#) - [metastore] The table description should have a fixed height in the table page
- [HUE-6772](#) - [metadata] Cache Navigator password reads from scripts
- [HUE-6773](#) - [core] Clicks on Logs, Dump config with middle click are not supported
- [HUE-6776](#) - [search] Js error when clicking on the first two layout of a new dashboard
- [HUE-6777](#) - [fb] Drag & drop upload not enabled if editor was loaded before
- [HUE-6778](#) - [frontend] Visual hint on click or d'n'd to import in Hue 4
- [HUE-6779](#) - [job] Disable click around checkbox in list of jobs
- [HUE-6780](#) - [s3] Correctly infer and display region when connected to S3 by endpoint
- [HUE-6781](#) - [job] Workflow duration time is not humanized

- [HUE-6783](#) - [home] Don't close the modals until all files marked for deletion are actually deleted
- [HUE-6785](#) - [doc2] Return more useful error message when users try to create a directory within a read-only shared directory
- [HUE-6786](#) - [metastore] DB configuration list shouldn't overflow into table list
- [HUE-6787](#) - [oozie] Hide calendar on modal close for any submit action
- [HUE-6789](#) - [beeswax] Fix alignment of the assist headers
- [HUE-6790](#) - [assist] Improve the assist functions filter
- [HUE-6791](#) - [search] Protect against pivot facets conflicting with nested facets
- [HUE-6792](#) - [core] Middle/right click is not working as a link on main blue action button
- [HUE-6794](#) - [assist] Show columns in the right assistant
- [HUE-6795](#) - [fb] Fix minor logic error in trash directory creation
- [HUE-6796](#) - [indexer] Faster installations of the examples
- [HUE-6797](#) - [dashboard] Grid widget operation icons are not grayed much
- [HUE-6798](#) - [frontend] Rename main action button depending on the selected app
- [HUE-6800](#) - [frontend] Re-word the welcome tour
- [HUE-6801](#) - [frontend] Set a minimum height to the log resizer binding
- [HUE-6802](#) - [editor] Don't check for column existence in the location handler when a column is prefixed with a table name or alias
- [HUE-6803](#) - [assist] Only refresh the assist cache for DDL statements that affect the assist structure
- [HUE-6806](#) - [frontend] Avoid double init of Nicescroll on the context popover
- [HUE-6807](#) - [search] The rolling time picker should consider the user's timezone
- [HUE-6808](#) - [frontend] Replace external links with _blank target to open the sample popup instead
- [HUE-6809](#) - [home] Set a min width for the name column in the document list
- [HUE-6810](#) - [assist] Improve positioning and visibility of side panel toggles
- [HUE-6811](#) - [editor] Prevent displaying small black lines in the upper left corner on editor load
- [HUE-6812](#) - [frontend] Swap sortable column style from top to bottom blue border
- [HUE-6813](#) - [metastore] Fix database comment last character truncation
- [HUE-6814](#) - [search] Only return distinct usernames in security impersonate dropdown
- [HUE-6815](#) - [frontend] Improve the user experience for the top search bar
- [HUE-6816](#) - [frontend] Fix flexbox sizing of the popover sample tab on Safari
- [HUE-6817](#) - [editor] Do a better column name match for the scroll to column functionality
- [HUE-6819](#) - [oozie] Set generic widget for generic actions in oozie graph
- [HUE-6820](#) - [frontend] Catch all the JS errors and send them to the backend
- [HUE-6821](#) - [home] Add star to favorite on Hue 4
- [HUE-6823](#) - [frontend] Sending JS errors to the backend should be transparent and still show the stacktrace on the console
- [HUE-6824](#) - [fb] Fix possible JS error after permanently deleting files
- [HUE-6825](#) - [autocomplete] Prevent exception when there's no identifierChain given for a table
- [HUE-6826](#) - [oozie] Better slicing of action commands property
- [HUE-6827](#) - [editor] Fix padding on download modal body
- [HUE-6828](#) - [editor] Fix vertical alignment of save query result file chooser button
- [HUE-6829](#) - [editor] Set download dialog overwrite default value to false to improve UX
- [HUE-6830](#) - [editor] Introduce copy to clipboard for results
- [HUE-6831](#) - [metastore] Editing view comments should use ALTER VIEW instead of TABLE
- [HUE-6833](#) - [editor] Disable query builder tab for non-sql editors
- [HUE-6835](#) - [autocomplete] Update the Impala auto-completer to match the latest Impala version
- [HUE-6836](#) - [jib] Reset select all checkboxes on interface change
- [HUE-6837](#) - [frontend] Fix alignment of hueCheckbox and hueCheckAll icons
- [HUE-6839](#) - [home] Avoid action icons to be shown unstyled at page load
- [HUE-6840](#) - [notebook] Workaround lack of status flag in Impala profile page

- [HUE-6841](#) - [oozie] Copy examples to HDFS as the hue user
- [HUE-6842](#) - [core] Correctly append the list of custom apps
- [HUE-6843](#) - [metastore] Show partition icon on table browser sends to Hue 3
- [HUE-6844](#) - [frontend] Make all custom apps embeddable
- [HUE-6845](#) - [frontend] Fix custom app mako generation
- [HUE-6846](#) - [frontend] Make all proxy apps embeddable
- [HUE-6847](#) - [metastore] Do not display if table is compressed as Hive information is incorrect
- [HUE-6848](#) - [job] Add doAs impersonation when fetching job logs
- [HUE-6850](#) - [metastore] Rely on default SQL editor type for querying data
- [HUE-6852](#) - [editor] Scroll left does not disappear when going to metastore from editor
- [HUE-6853](#) - [job] On mini job clicking on breadcrumbs shouldn't close the popup
- [HUE-6854](#) - [home] Clicking on a folder should open it
- [HUE-6855](#) - [frontend] Set a default application tooltip sometimes doesn't disappear
- [HUE-6856](#) - [search] Protect against reflected XSS in search query parameters
- [HUE-6857](#) - [metadata] Add skeleton of troubleshooting API
- [HUE-6858](#) - [indexer] Add load data action
- [HUE-6859](#) - [index] Support creating collection aliases without page refresh
- [HUE-6861](#) - [search] Do not 500 when Search is not enabled
- [HUE-6862](#) - [job] Workflow filtering on username does not happen
- [HUE-6863](#) - [job] Kill buttons are enabled even when the user is not allowed
- [HUE-6864](#) - [job] Check for permission before killing a job
- [HUE-6866](#) - [core] Js error popup does not re-appear if same as previous one which was closed
- [HUE-6868](#) - [search] Field analysis popup shows up too high if down scrolling
- [HUE-6870](#) - [job] Improve end user experience with bubbling up of common actions and properties
- [HUE-6874](#) - [search] Stabilize some nested facets
- [HUE-6877](#) - [rdbms] Server error label misspelling
- [HUE-6880](#) - [fb] Missing bottom gray border table header styling
- [HUE-6881](#) - [editor] Display table type in the SQL context popover
- [HUE-6882](#) - [metastore] Fix issues with the description overflow toggle
- [HUE-6883](#) - [metastore] Allow deletion of table comments
- [HUE-6884](#) - [core] Allow to blacklist the Hive
- [HUE-6885](#) - [frontend] Catch all http 502 and show an HTML stripped version of the message
- [HUE-6886](#) - [oozie] SLA variable in workflow action are not getting retrieved
- [HUE-6888](#) - [oozie] SLA page is broken in Hue 4
- [HUE-6889](#) - [oozie] Force wrap oozie IDs on the detail page
- [HUE-6890](#) - [autocomplete] Add support for database prefixed UDFs
- [HUE-6891](#) - [importer] Incorrect Warning for tables starting with #
- [HUE-6894](#) - [core] Remove cluster config call on login page
- [HUE-6895](#) - [frontend] Disable errorcatcher when in Django debug mode
- [HUE-6896](#) - [frontend] Errorcatcher should not log jQuery triggered events
- [HUE-6897](#) - [oozie] Ensure that examples_dir exists before installing examples
- [HUE-6898](#) - [oozie] Editor doesn't update job status to "finished" for finished Pig jobs
- [HUE-6899](#) - [core] Allow 403 to be accessed by regular users
- [HUE-6900](#) - [useradmin] Do not link app title to list of users if user is not a super user
- [HUE-6901](#) - [oozie] Editor is missing certain regular actions
- [HUE-6903](#) - [dataeng] Backport several useful Jiras
- [HUE-6905](#) - [frontend] Minor user interface bugs in old Pig Editor and Job Designer
- [HUE-6906](#) - [frontend] 'Add more' should open in a new tab and not throw a 404 in Hue 4
- [HUE-6907](#) - [search] The dashboard header is bigger than the other headers
- [HUE-6908](#) - [job] Add optional Spark session list

- [HUE-6909](#) - [libsentry] Use the new sentry config property `sentry.service.client.server.rpc-addresses`
- [HUE-6910](#) - [indexes] Display dynamic and copy fields
- [HUE-6911](#) - [indexes] Display sample data of schemaless collections
- [HUE-6912](#) - [sqoop] Securing Database password in oozie
- [HUE-6922](#) - [fb] When typing in the search box, last 1 or few characters get deleted when the list is refreshed
- [HUE-6923](#) - [fb] Fix table header borders
- [HUE-6925](#) - [indexes] Enable column filtering
- [HUE-6926](#) - [indexes] Refresh out-of-sync collections
- [HUE-6928](#) - [assist] Add filter to right assistant
- [HUE-6930](#) - [importer] Libs filechooser input should be shorter
- [HUE-6931](#) - [indexes] Hide table headers when tables are empty
- [HUE-6932](#) - [indexes] Add field sample popover
- [HUE-6933](#) - [indexes] Show a message when there's no sample data
- [HUE-6934](#) - [importer] Taking the Sqoop jar paths from UI parameters/asking from end-user
- [HUE-6936](#) - [indexer] Move SQL indexing into its own flag
- [HUE-6937](#) - [importer] Improve UX by not hiding the pre-filled wizard destination
- [HUE-6938](#) - [importer] Restyle section headers
- [HUE-6939](#) - [editor] Improve interactions with the risk indicator
- [HUE-6940](#) - [editor] Cancel risk requests if editing while risk check is running
- [HUE-6941](#) - [editor] Don't show a pointer cursor for the risk check spinner
- [HUE-6942](#) - [indexer] 404 error when creating an index from a file that has a `_version_` field
- [HUE-6943](#) - [indexer] Error when input file does not have any headers
- [HUE-6944](#) - [indexer] Support manual creation of collection
- [HUE-6945](#) - [search] Aliases cannot be retrieved with Solr 4
- [HUE-6948](#) - [indexer] Warn when indexer name is invalid
- [HUE-6949](#) - [indexer] Cleaner error message when coming from Solr
- [HUE-6950](#) - [sam] Create home directory for new user login
- [HUE-6951](#) - [indexes] Index with no data has an empty field section
- [HUE-6952](#) - [doc] Remove warning description
- [HUE-6953](#) - [indexer] Support skipping field in non task mode
- [HUE-6954](#) - [indexer] Support different number of shards, replication at index creation
- [HUE-6955](#) - [metadata] Extend query suggest caching to 1 week
- [HUE-6957](#) - [importer] Primary Key dropdown opens on selecting Default field dropdown
- [HUE-6959](#) - [assist] Fix collection opening from left assist
- [HUE-6960](#) - [frontend] Only suggest facet values and no results when autocompleting facets in the top search
- [HUE-6961](#) - [frontend] Filechooser inputs sometimes have rounder corners on the right side of the input
- [HUE-6962](#) - [importer] Index field operations overlap with the submit button
- [HUE-6965](#) - [importer] Supporting import in different file types with different delimiters
- [HUE-6968](#) - [metadata] Gather table stats via Impala
- [HUE-6969](#) - [job] Start time, Finish time, Elapsed time of tasks of MR jobs are not formatted
- [HUE-6971](#) - [job] Schedule list is filtered by default on 7 days which hides most of them
- [HUE-6972](#) - [metadata] Sometimes the DB owner is shown as ()
- [HUE-6973](#) - [indexer] Open up the index page after creating a new index
- [HUE-6974](#) - [importer] Refresh Assist collection list after create table wizard success
- [HUE-6975](#) - [core] 500 in configuration on `/hue/desktop/dump_config`
- [HUE-6976](#) - [editor] Suggest fix for missing filters risk
- [HUE-6977](#) - [editor] Fix issue where the autocomplete dropdown doesn't show for column completion
- [HUE-6979](#) - [editor] Sometimes editor title/summary are truncated too much
- [HUE-6981](#) - [oozie] Configuration text filter spacing is a bit off
- [HUE-6984](#) - [editor] Use TeraGen instead of TeraSort

- [HUE-6988](#) - [job] After killing the MR job, the interface hangs
- [HUE-6989](#) - [job] Clicking on the jobs button doesn't allow me to go back to the editor with the back button of the browser
- [HUE-6990](#) - [editor] The syntax checker should be aware of table aliases
- [HUE-6992](#) - [editor] Unable to paste via right click menu
- [HUE-6993](#) - [sqoop] Support for configurable number of mappers
- [HUE-6995](#) - [oozie] Set minimum width of a workflow node
- [HUE-6996](#) - [frontend] Clicking on home doesn't allow to go back in history
- [HUE-6997](#) - [core] Leaflet_tile_layer should allow case sensitive URIs
- [HUE-6999](#) - [core] Deleting admin user message yes button keeps circling
- [HUE-7000](#) - [assist] HBase assist does not load the table the second time
- [HUE-7001](#) - [assist] Remember the last selected HBase cluster
- [HUE-7002](#) - [core] Welcome to Hue 4 tour could be also skipped by clicking on the white background
- [HUE-7004](#) - [notebook] Loading query with empty session param shows several JS errors
- [HUE-7005](#) - [oozie] sharing coordinator or bundle immediately after creating fails
- [HUE-7006](#) - [core] Ability in HUE to change Cookie name from default "sessionid"
- [HUE-7008](#) - [autocomplete] Minify the generated parsers
- [HUE-7009](#) - [metadata] Automatic upload of table stats if missing
- [HUE-7010](#) - [sqoop] Take configured databases from librdms instead of hardcoding the list of drivers
- [HUE-7011](#) - [job] JS error on workflow with tasks without a startTime
- [HUE-7012](#) - [job] Wrong spelling of Job has not tasks
- [HUE-7013](#) - [backend] Provide Kerberos Mutual Authentication disable switch in "hue.ini"
- [HUE-7015](#) - [frontend] Set default background color to white
- [HUE-7016](#) - [libsentry] HA retry on SOLR component does not set the component value
- [HUE-7017](#) - [importer] Add JDBC option in both custom and config modes
- [HUE-7018](#) - [job] disable_killing_jobs option is not supported in the new job browser
- [HUE-7019](#) - [fb] show_download_button is not integrated in the backend
- [HUE-7020](#) - [indexer] Provide the real lib_path at the Morphline config generation
- [HUE-7022](#) - [assist] Enable listing the Solr collection fields
- [HUE-7023](#) - [assist] Add a create button to the HDFS panel
- [HUE-7024](#) - [assist] The current selected item should have an 'open' action too
- [HUE-7025](#) - [editor] Move copy to clipboard outside the download dropdown
- [HUE-7026](#) - [editor] Add an HBase connector
- [HUE-7027](#) - [fb] Opening a file directly via editing path lands on Hue 3
- [HUE-7028](#) - [editor] Don't mark where and limit clauses on hover
- [HUE-7029](#) - [job] Switching between log types should reset the log content before loading the new one
- [HUE-7030](#) - [job] Style Oozie SLA not available page
- [HUE-7031](#) - [oozie] Promote tabs to the usual Documents icon for browsing documents
- [HUE-7033](#) - [job] Switching from Hue 3 to Hue 4 should remap the old URL to the new one
- [HUE-7035](#) - [frontend] Close editor sessions panel on app switch
- [HUE-7036](#) - [fb] Disable full page refresh on folder open
- [HUE-7039](#) - [oozie] Remove JS auto center of the workflow actions
- [HUE-7040](#) - [oozie] Pressing Esc on share popup throws JS error in Hue4
- [HUE-7041](#) - [fb] Items in menu come and go when doing a right clicking on a file
- [HUE-7043](#) - [frontend] Switch fa-tags with fa-list for list of files links in the apps
- [HUE-7044](#) - [oozie] Be consistent with the Schedule name in Hue 4
- [HUE-7045](#) - [hive] Avoid use database call for SET queries
- [HUE-7046](#) - [frontend] Fix open.link on Hue 3
- [HUE-7048](#) - [oozie] User prompting doesn't work in share document popup in hue4
- [HUE-7049](#) - [assist] Extract a generic context popover component

- [HUE-7050](#) - [assist] Add a feature flag for the new generic context popover
- [HUE-7051](#) - [importer] Add SQL editor to the importer for RDBMS
- [HUE-7052](#) - [desktop] Retain last_modified when sharing sample docs in sync_documents
- [HUE-7053](#) - [search] Indexer requires options to avoid jar classes conflicts
- [HUE-7055](#) - [importer] Option to add column other than primary key to split-by
- [HUE-7056](#) - [importer] Minimal arrow linking step 1 to 2
- [HUE-7057](#) - [search] Make widget dimensions draggable
- [HUE-7060](#) - [search] Merge initial layout choices to one single widget and two columns grid
- [HUE-7061](#) - [search] Add the possibility of add/remove the left layout column
- [HUE-7064](#) - [editor] Integrate sharing button into the Editor
- [HUE-7065](#) - [dashboard] Integrate sharing button into the Dashboard
- [HUE-7066](#) - [frontend] Missing grunt uglify dependency on packages.json
- [HUE-7067](#) - [dashboard] Refresh facet result when reordering sub-facets
- [HUE-7068](#) - [frontend] Update knockout sortable binding to the latest version
- [HUE-7069](#) - [editor] Remove the spinner for risk checking
- [HUE-7070](#) - [importer] Prefill the destination name for the database importer
- [HUE-7071](#) - [frontend] The left nav should always be on top
- [HUE-7072](#) - [importer] Disable fields that are not marked as keep
- [HUE-7073](#) - [backend] When user is logged in using LdapBackEnd, restrict certain characters from the Hue login name.
- [HUE-7074](#) - [frontend] Extend hueDebug to accept the id of the element as well
- [HUE-7075](#) - [assist] Allow to refresh other types than Hive and Impala
- [HUE-7079](#) - [editor] Keep the last 50 risk responses and clear the risks on statement change
- [HUE-7080](#) - [oozie] Oozie HiveServer2 action is prompting for and 'Impalad hostname'
- [HUE-7082](#) - [aws] Create bucket fails with error "Failed to create S3 bucket 'xxxx':Bad Request"
- [HUE-7084](#) - [useradmin] Remove @domain part of username when creating home directory for LDAP user
- [HUE-7085](#) - [useradmin] Remove unused view_user view
- [HUE-7086](#) - [core] Clicking on a task item to select it closes the task popup
- [HUE-7087](#) - [frontend] Skip showing assist.db.refresh output links on the history panel
- [HUE-7090](#) - [security] Protect from missing groupName in the list_sentry_roles_by_group API
- [HUE-7091](#) - [assist] The assistant should suggest adding a filter for missing partition columns in filter
- [HUE-7092](#) - [autocomplete] Adjust the parser to also report location of select lists
- [HUE-7093](#) - [frontend] Unify table headers and sorters styles across the various apps
- [HUE-7094](#) - [frontend] Filechooser should never navigate by default to an empty path
- [HUE-7095](#) - Table import not working
- [HUE-7096](#) - [editor] Presentation mode revamp
- [HUE-7097](#) - [notebook] Remove export to Jupyter
- [HUE-7098](#) - [presentation] Toggle hiding/showing the code snippets
- [HUE-7099](#) - [notebook] Display variables on top
- [HUE-7100](#) - [editor] Switch ko model between editor and notebook presentation mode
- [HUE-7104](#) - [core] Hue and Isilon issue "UnsupportedOperationException: Unknown op 'GETTRASHROOT'"
- [HUE-7105](#) - [importer] Column list is not aligned on large screens
- [HUE-7106](#) - [oozie] New oozie doc url doesn't update after saving the page
- [HUE-7107](#) - [core] Hue 4 group access doesn't reflect on Interface
- [HUE-7108](#) - [desktop] Add config check to see if the migration history table is up to date
- [HUE-7109](#) - [useradmin] Add/Sync LDAP Group should refresh the group list and show a sync message
- [HUE-7110](#) - [editor] Focus on the SQL editor when opening Hue
- [HUE-7112](#) - [oozie] Do not escape characters in decision node condition
- [HUE-7114](#) - [filebrowser] Edit File could be in a fixed width font
- [HUE-7115](#) - [autocomplete] Show popular joins above popular filters

- [HUE-7116](#) - [editor] Fix syntax checker js error when table ref is missing
- [HUE-7117](#) - [frontend] Harmonize all the monospace occurrences and move to Roboto Mono
- [HUE-7118](#) - [editor] Take statement location into account when marking query compatibility warnings
- [HUE-7120](#) - [presentation] Set a minimum height for the editors
- [HUE-7122](#) - [autocomplete] Don't add db prefix in popular join conditions for Hive
- [HUE-7123](#) - [autocomplete] Include the join type if given when suggesting popular joins
- [HUE-7124](#) - [frontend] Use the originalName attribute for inline autocomplete in the global search
- [HUE-7125](#) - [importer] Dropdown shrinks in the 'external database' section
- [HUE-7127](#) - [backend] Fix for Python requests Kerberos/GSSAPI authentication library fails to authenticate kerberos requests to the same destination host.
- [HUE-7128](#) - [core] Create an ENABLE_DOWNLOAD property in 'desktop' conf and use it for the editor and search dashboards
- [HUE-7131](#) - [importer] Revert HUE-7051
- [HUE-7132](#) - [frontend] Fix Selectize default placeholder i18n
- [HUE-7133](#) - [oozie] Bulk add parameters parsed from XML
- [HUE-7134](#) - [oozie] Allow vertical resize of Sqoop arguments
- [HUE-7135](#) - [presentation] Add SQL comment as snippet title or text widgets
- [HUE-7136](#) - [frontend] Remove references to download.js
- [HUE-7137](#) - [editor] Move copy to clipboard inside the download dropdown
- [HUE-7138](#) - [editor] Open notebook should work also when the editor has not been loaded before
- [HUE-7139](#) - [conf] Remove non-standard double quotes from the configuration files
- [HUE-7141](#) - [dashboard] Menu shows an incomplete dropdown
- [HUE-7144](#) - [core] Switch document list icon in editors to the document one
- [HUE-7148](#) - [frontend] Set a default size for SVG icons inside buttons
- [HUE-7149](#) - [presentation] Missing common share import in Hue 3
- [HUE-7150](#) - [presentation] Skip storing the left and right assist status when toggling presentation mode
- [HUE-7151](#) - [editor] Refactor previousChartOptions to avoid duplication
- [HUE-7152](#) - [editor] Enable execute query shortcut inside a variable input
- [HUE-7153](#) - [core] Analytics call to add logging in the backend
- [HUE-7154](#) - [hive] Replace NoSuchObjectException by QueryServerException
- [HUE-7155](#) - [backend] Hue needs to support TLSv1_1 and TLSv1_2 via Thrift connections
- [HUE-7157](#) - [presentation] Align the top variables better
- [HUE-7158](#) - [editor] Remove border to the grid results column list
- [HUE-7159](#) - [fb] Remove history dropdown
- [HUE-7160](#) - [importer] Make arrow in between steps bigger
- [HUE-7162](#) - [assist] Make the document type selection scrollable
- [HUE-7165](#) - [aws] hue s3 browser should provide accessibility for all buckets in v2 regions
- [HUE-7166](#) - [aws] deleting s3 bucket throws console error
- [HUE-7167](#) - [search] Enable swapping dashboard columns
- [HUE-7168](#) - [editor] Add keyboard shortcut for presentation mode and make esc work all the time
- [HUE-7169](#) - [editor] Force re-render when swapping between chart and grid mode
- [HUE-7170](#) - [indexes] Analysis popup hidden a bit on non refreshed collection page
- [HUE-7172](#) - [indexes] Analysis popup should disappear when navigating to the list of indexes
- [HUE-7173](#) - [indexes] Add name label to the create alias modal
- [HUE-7175](#) - [fb] Right click on file, click on a disable action actually clicks on the background
- [HUE-7176](#) - [metastore] Open in metastore from Assist does not select the first tab in Hue 3
- [HUE-7177](#) - [assist] Right click "Open in Importer"
- [HUE-7179](#) - [core] Avoid issue with searching documents on assist or global search box
- [HUE-7180](#) - [search] Hide horizontal column swap when there is only one
- [HUE-7181](#) - [assist] Collection field list does not get refreshed

- [HUE-7183](#) - [indexer] Properly bubble up errors when index creation fails
- [HUE-7184](#) - [search] Protect against nested facets with empty buckets
- [HUE-7185](#) - [indexer] Go back to the index list after deleting one index
- [HUE-7187](#) - [editor] Save button should be enabled once table name being filled
- [HUE-7188](#) - [editor] The save dropdown should be under the horizontal scrollbar
- [HUE-7189](#) - [editor] Query result modal is shown only once
- [HUE-7192](#) - [desktop] Hue Search is failing with DatabaseError field_cast_sql deprecated after Django 1.8
- [HUE-7193](#) - [core] /desktop/debug/is_alive should allow GET requests
- [HUE-7194](#) - [assist] Fix broken documents panel refresh
- [HUE-7196](#) - [assist] Fix JS error for missing Object.values in IE 11
- [HUE-7197](#) - [editor] Fix IE 11 JS errors in syntax checker
- [HUE-7198](#) - [editor] Remove unused onBlur code to improve IE 11 debug performance
- [HUE-7200](#) - [metastore] Refresh table and database js errors
- [HUE-7201](#) - [importer] 'Config' object has no attribute 'get' when Solr is not installed
- [HUE-7205](#) - [editor] Fix leaking interval in logScroller
- [HUE-7206](#) - [editor] Fix stack overflow from fileChooser binding in IE 11
- [HUE-7207](#) - [frontend] Fix fileChooser search for IE 11
- [HUE-7209](#) - [editor] Make sure pubsub subscribes continue to work when switching between editor types
- [HUE-7210](#) - [metastore] Fix the metastore refresh button
- [HUE-7211](#) - [metastore] Invalidate the Impala cache when pressing refresh in the metastore
- [HUE-7212](#) - [editor] Properly scroll to the executing query when selected queries are executed
- [HUE-7213](#) - [editor] Mark the correct statement when executing selection after the first line
- [HUE-7216](#) - [presentation] Close the editor mode when leaving by other links than the exit button
- [HUE-7218](#) - [editor] Offer links if row values are pure http links in the record popup
- [HUE-7219](#) - [dashboard] New columns are one slot too small when adding a new one
- [HUE-7223](#) - [autocomplete] Fix issue where the old autocompleter inserts extra characters on exception
- [HUE-7225](#) - [autocomplete] Prevent new line on enter when a suggestion equals the typed text
- [HUE-7226](#) - [frontend] Add navigation logging
- [HUE-7227](#) - [frontend] Log usage of popular values in the autocomplete
- [HUE-7231](#) - [assist] Keep track of the last opened documents folder
- [HUE-7234](#) - [editor] New variable should be empty
- [HUE-7240](#) - [impala] Workaround non support of partition_spec in DESCRIBE FORMATTED
- [HUE-7241](#) - [assist] Only refresh the docs assist when saving new documents in the editor
- [HUE-7242](#) - [frontend] Properly dispose the context popover contents for the global search results
- [HUE-7243](#) - [assist] Make sure the SQL assist panel is opened from the context popover show in assist action
- [IMPALA-992](#) - Rerun past queries from history in shell
- [IMPALA-1470](#) - Fix error message with catalog down
- [IMPALA-1478](#) - Improve error message when subquery is used in the ON clause
- [IMPALA-1882](#) - Remove ORDER BY restriction from first_value()/last_value()
- [IMPALA-1891](#) - Statestore won't send deletions in initial non-delta topic
- [IMPALA-2373](#) - Extrapolate row counts for HDFS tables.
- [IMPALA-2525](#) - Treat parquet ENUMs as STRINGs when creating impala tables.
- [IMPALA-2615](#) - support [[nodiscard]] on Status
- [IMPALA-3040](#) - addendum: use specific_build_type_timeout for slow builds
- [IMPALA-3208](#) - max_row_size option
- [IMPALA-3504](#) - UDF for current timestamp in UTC
- [IMPALA-3894](#) - Change the behavior parsing date "YY"
- [IMPALA-3905](#) - HdfsScanner::GetNext() for Avro, RC, and Seq scans.
- [IMPALA-3931](#) - arbitrary fixed-size uda intermediate types
- [IMPALA-3937](#) - Deprecate --be_service_threads

- [IMPALA-4039](#) - Increase width of Operator column in runtime profile
- [IMPALA-4086](#) - Add benchmark for simple scheduler
- [IMPALA-4164](#) - Avoid overly aggressive inlining in LLVM IR
- [IMPALA-4192](#) - Disentangle Expr and ExprContext
- [IMPALA-4276](#) - Profile displays non-default query options set by planner
- [IMPALA-4418](#) - Fixes extra blank lines in query result
- [IMPALA-4622](#) - Add ALTER COLUMN statement.
- [IMPALA-4623](#) - Enable file handle cache
- [IMPALA-4669](#) - Add Kudu's RPC, util, and security libraries
- [IMPALA-4674](#) - Port spilling ExecNodes to new buffer pool
- [IMPALA-4687](#) - Get Impala working against HBase 2.0
- [IMPALA-4703](#) - reservation denial debug action
- [IMPALA-4737](#) - Prevent SIGUSR1 from killing daemons when minidumps are disabled
- [IMPALA-4794](#) - Grouping distinct agg plan robust to data skew
- [IMPALA-4795](#) - Allow fetching function obj from catalog using signature
- [IMPALA-4826](#) - Fix error during a scan on repeated root schema in Parquet.
- [IMPALA-4833](#) - Use scheduling information to make per-node memory reservation tight
- [IMPALA-4861](#) - READ_WRITE warning on CREATE TABLE LIKE PARQUET.
- [IMPALA-4862](#) - make resource profile consistent with backend behaviour
- [IMPALA-4866](#) - Hash join node does not apply limits correctly
- [IMPALA-4905](#) - Don't send empty insert status to coordinator
- [IMPALA-4925](#) - Cancel finstance if query has finished
- [IMPALA-5016](#) - Simplify COALESCE() in SimplifyConditionalsRule.
- [IMPALA-5031](#) - UBSAN clean and method for testing UBSAN cleanliness
- [IMPALA-5036](#) - Parquet count star optimization
- [IMPALA-5061](#) - Populate null_count in parquet::statistics
- [IMPALA-5085](#) - large rows in BufferedTupleStreamV2
- [IMPALA-5104](#) - Admit queries with mem equal to proc mem_limit
- [IMPALA-5108](#) - idle_session_timeout kicks in later than expected
- [IMPALA-5109](#) - Increase range of backend latency histogram
- [IMPALA-5116](#) - Remove deprecated hash_* types in gutil
- [IMPALA-5158](#) - Account for difference between process memory consumption and memory used by queries
- [IMPALA-5160](#) - adjust spill buffer size based on planner estimates
- [IMPALA-5167](#) - Reduce the number of Kudu clients created
- [IMPALA-5240](#) - Allow config of number of disk I/O threads per disk type
- [IMPALA-5275](#) - Avoid printing status stack trace on hot paths
- [IMPALA-5280](#) - Coalesce chains of OR conditions to an IN predicate
- [IMPALA-5286](#) - Query with Kudu col name w/ different casing from 'order by' fails
- [IMPALA-5316](#) - Adds last_day() function
- [IMPALA-5327](#) - Handle return of JNI GetStringUTFChar
- [IMPALA-5336](#) - Fix partition pruning when column is cast
- [IMPALA-5347](#) - reduce codegen overhead of timestamp trunc()
- [IMPALA-5350](#) - Tidy up thread groups for first exec threads
- [IMPALA-5352](#) - Age out unused file handles from the cache
- [IMPALA-5354](#) - INSERT hints for Kudu tables
- [IMPALA-5376](#) - Loads all TPC-DS tables
- [IMPALA-5386](#) - Fix ReopenCachedHdfsFileHandle failure case
- [IMPALA-5389](#) - simplify BufferDescriptor lifetime
- [IMPALA-5407](#) - Fix crash in HdfsSequenceTableWriter
- [IMPALA-5412](#) - Fix scan result with partitions on same file

- [IMPALA-5420](#) - Skip ACL fetch if the acl bit is not set.
- [IMPALA-5427](#) - Fix race between CRS::UpdateQueryStatus() and beeswax RPCs
- [IMPALA-5431](#) - Remove redundant path exists checks during table load
- [IMPALA-5433](#) - Mark single-argument Status c'tors as explicit
- [IMPALA-5446](#) - dropped Sorter::Reset() status
- [IMPALA-5477](#) - Fix minidump-2-core tool
- [IMPALA-5479](#) - Propagate the argument type in RawValue::Compare()
- [IMPALA-5480](#) - Improve missing filters message
- [IMPALA-5481](#) - Clarify RowDescriptor ownership
- [IMPALA-5482](#) - fix git checkout when workloads are modified
- [IMPALA-5483](#) - Automatically disable codegen for small queries
- [IMPALA-5484](#) - Fix LICENSE issues discovered by IPMC in 2.9 vote
- [IMPALA-5488](#) - Fix handling of exclusive HDFS file handles
- [IMPALA-5489](#) - Improve Sentry authorization for Kudu tables
- [IMPALA-5492](#) - Fix incorrect newline character in the LDAP message
- [IMPALA-5494](#) - Fixes the selectivity of NOT IN predicates
- [IMPALA-5495](#) - Improve error message if no impalad role is configured
- [IMPALA-5497](#) - spilling hash joins that output build rows hit OOM
- [IMPALA-5498](#) - Support for partial sorts in Kudu INSERTs
- [IMPALA-5499](#) - avoid ephemeral port conflicts
- [IMPALA-5500](#) - Reduce catalog update topic size
- [IMPALA-5504](#) - Fix TupleIsNullPredicate evaluation.
- [IMPALA-5506](#) - Add stdin description to help information of query_file option
- [IMPALA-5507](#) - Add clear description to help information of KEYVAL option
- [IMPALA-5511](#) - Add process start time to debug web page
- [IMPALA-5513](#) - Fix display message exception when using invalid KEYVAL
- [IMPALA-5514](#) - Throw an error when --ldap_password_cmd is used without LDAP auth
- [IMPALA-5517](#) - Allow IMPALA_LOGS_DIR to be overridden
- [IMPALA-5520](#) - TopN node periodically reclaims old allocations
- [IMPALA-5524](#) - Fixes NPE during planning with DISABLE_UNSAFE_SPILLS=1
- [IMPALA-5529](#) - Add additional function signatures for TRUNC()
- [IMPALA-5531](#) - Fix correctness issue in correlated aggregate subqueries
- [IMPALA-5532](#) - Stack-allocate compressors in RowBatch (de)serialization
- [IMPALA-5536](#) - Fix TCLIService thrift compilation on Hive 2
- [IMPALA-5539](#) - Fix Kudu timestamp with -use_local_tz_for_unix_ts
- [IMPALA-5540](#) - Revert Sentry version back to 5.13
- [IMPALA-5546](#) - Allow creating unpartitioned Kudu tables
- [IMPALA-5547](#) - Rework FK/PK join detection.
- [IMPALA-5548](#) - Fix some minor nits with HDFS parquet column readers
- [IMPALA-5549](#) - Remove deprecated fields from CatalogService API
- [IMPALA-5551](#) - Fix AggregationNode::Close() when Prepare() fails
- [IMPALA-5554](#) - sorter DCHECK on null column
- [IMPALA-5560](#) - always store CHAR(N) inline in tuple
- [IMPALA-5567](#) - race in fragment instance teardown
- [IMPALA-5570](#) - fix spilling null-aware anti join
- [IMPALA-5572](#) - Timestamp codegen for text scanner
- [IMPALA-5573](#) - Add decimal codegen in text scanner
- [IMPALA-5579](#) - Fix IndexOutOfBoundsException in GetTables metadata request
- [IMPALA-5580](#) - fix Java UDFs that return NULL strings
- [IMPALA-5582](#) - Store sentry privileges in lower case

- [IMPALA-5586](#) - Null-aware anti-join can take a long time to cancel
- [IMPALA-5588](#) - Reduce the frequency of fault injection
- [IMPALA-5591](#) - set should handle negative values
- [IMPALA-5594](#) - don't import shaded classes
- [IMPALA-5595](#) - Only set KuduScanner timestamp feature flag if necessary
- [IMPALA-5598](#) - Fix excessive dumping in MemLimitExceeded
- [IMPALA-5602](#) - Fix query optimization for kudu and datasource tables
- [IMPALA-5611](#) - KuduPartitionExpr holds onto memory unnecessarily
- [IMPALA-5612](#) - join inversion should factor in parallelism
- [IMPALA-5615](#) - Fix compute incremental stats for general partition exprs
- [IMPALA-5616](#) - Add --enable_minidumps startup flag
- [IMPALA-5617](#) - Include full workload name in tpch_nested query filenames
- [IMPALA-5618](#) - buffered-tuple-stream-v2 fixes
- [IMPALA-5623](#) - Fix lag() on STRING cols to release UDF mem
- [IMPALA-5627](#) - fix dropped statuses in HDFS writers
- [IMPALA-5629](#) - avoid expensive list::size() call
- [IMPALA-5630](#) - Add Kudu client version as a common metric
- [IMPALA-5636](#) - Change the metadata in parquet
- [IMPALA-5638](#) - Fix Kudu table set tblproperties inconsistencies
- [IMPALA-5641](#) - mem-estimate should never be less than mem-reservation
- [IMPALA-5643](#) - Add total number of threads created per group to /threadz
- [IMPALA-5644](#) - Fail queries early when their minimum reservation is too high to execute within the given mem_limit
- [IMPALA-5648](#) - fix count(*) mem estimate regression
- [IMPALA-5650](#) - Make sum_init_zero a SUM function
- [IMPALA-5652](#) - deprecate unlimited process mem_limit
- [IMPALA-5657](#) - FunctionCallExpr.toSql() and clone() ignore 'IGNORE NULLS' case
- [IMPALA-5658](#) - addtl. process/system-wide memory metrics
- [IMPALA-5659](#) - Begin standardizing treatment of thirdparty libraries
- [IMPALA-5659](#) - glog/gflags should be dynamically linked
- [IMPALA-5661](#) - buffer pool limit
- [IMPALA-5666](#) - ASAN poisoning for MemPool and BufferPool
- [IMPALA-5670](#) - Misc. tidying of ExecEnv
- [IMPALA-5676](#) - avoid expensive consistency checks in BTSv2
- [IMPALA-5677](#) - limit clean page memory consumption
- [IMPALA-5679](#) - Fix Parquet count(*) with group by string
- [IMPALA-5681](#) - release reservation from blocking operators
- [IMPALA-5686](#) - Update a mini cluster Sentry property
- [IMPALA-5689](#) - Avoid inverting non-equi left joins
- [IMPALA-5691](#) - recalibrate mem limit for Q18
- [IMPALA-5696](#) - Enable cipher configuration when using TLS / Thrift
- [IMPALA-5709](#) - Remove mini-impala-cluster
- [IMPALA-5713](#) - always reserve memory for preaggs
- [IMPALA-5714](#) - Add OpenSSL to bootstrap_toolchain.py
- [IMPALA-5715](#) - (mitigation only) defer destruction of MemTrackers
- [IMPALA-5716](#) - Don't delete cmake_modules/* when enabling distcc
- [IMPALA-5722](#) - Fix string to decimal cast
- [IMPALA-5725](#) - coalesce() with outer join incorrectly rewritten
- [IMPALA-5739](#) - Correctly handle sles12 SP2
- [IMPALA-5742](#) - De-allocate buffer in parquet-reader on exit
- [IMPALA-5743](#) - Support TLS version configuration for Thrift servers

- [IMPALA-5744](#) - Add 'use_krpc' flag and create DataStream interface
- [IMPALA-5745](#) - Bump Breakpad version
- [IMPALA-5749](#) - coordinator race hits DCHECK 'num_remaining_backends_ > 0'
- [IMPALA-5750](#) - Catch exceptions from boost thread creation
- [IMPALA-5756](#) - start memory maintenance thread after metric creation
- [IMPALA-5757](#) - Make tbl property toSql deterministic
- [IMPALA-5769](#) - Add periodic minidump cleanup
- [IMPALA-5773](#) - Correctly account for memory used in data stream receiver queue
- [IMPALA-5774](#) - Prevent FindInSet() from reading off end of string
- [IMPALA-5775](#) - Allow shell to support TLSv1, v1.1 and v1.2
- [IMPALA-5776](#) - HdfsTextScanner::WritePartialTuple() writes the varlen data to an incorrect memory pool
- [IMPALA-5778](#) - Clarify logging around and usage of read_size startup option
- [IMPALA-5784](#) - Separate planner and user set query options in profile
- [IMPALA-5787](#) - Dropped status in KuduTableSink::Send()
- [IMPALA-5788](#) - Fix agg node crash when grouping by nondeterministic exprs
- [IMPALA-5791](#) - Make bootstrap_development.sh survive apt-get failure
- [IMPALA-5796](#) - CTAS for Kudu fails with expr rewrite
- [IMPALA-5798](#) - ASAN use-after-poison in Parquet decoder
- [IMPALA-5799](#) - Kudu DML can crash if schema has changed
- [IMPALA-5800](#) - Configure Squeasel's cipher suite and TLS version
- [IMPALA-5811](#) - Add 'backends' tab to query details pages
- [IMPALA-5812](#) - Fix NPE when joining on empty const select
- [IMPALA-5815](#) - right outer join returns invalid memory
- [IMPALA-5819](#) - DCHECK in HdfsTextScanner::Close()
- [IMPALA-5823](#) - fix SET_DENY_RESERVATION_PROBABILITY
- [IMPALA-5825](#) - Catch exceptions thrown by TSSLocketFactory c'tor
- [IMPALA-5838](#) - Improve errors on AC buffer mem rejection
- [IMPALA-5840](#) - Don't write page-level statistics in Parquet files.
- [IMPALA-5849](#) - Remove compile-time checks for OpenSSL > 1.0.0
- [IMPALA-5850](#) - Cast sender partition exprs under unions.
- [IMPALA-5852](#) - improve MINIMUM_RESERVATION_UNAVAILABLE error
- [IMPALA-5853](#) - fix GetResultSetMetadata() error message for invalid query id
- [IMPALA-5855](#) - Preaggregation crashes - unable to initialize hash table
- [IMPALA-5857](#) - avoid invalid free of hedged read metrics
- [IMPALA-5867](#) - Fix bugs parsing 2-digit year
- [IMPALA-5871](#) - KuduPartitionExpr incorrectly handles its child types
- [IMPALA-5885](#) - free runtime filter allocations in Parquet
- [IMPALA-5888](#) - free other local allocations in Parquet
- [IMPALA-5890](#) - Abort queries if scanner hits IO errors
- [IMPALA-5891](#) - fix PeriodicCounterUpdater initialization
- [IMPALA-5892](#) - Allow reporting status independent of fragment instance
- [IMPALA-2615](#) - annotate Status with [[nodiscard]]
- [IMPALA-5412](#) - Fix scan result with partitions on same file
- [KITE-762](#) - HiveAbstractDatasetRepository now uses the first URI from hive.metastore.uris property
- [KITE-1155](#) - Deleting an already deleted empty path should not fail the job
- [KUDU-871](#) - Allow tombstoned tablets to vote
- [KUDU-1125](#) - (part 1) catalog_manager: try to avoid unnecessarily rewriting tablet info
- [KUDU-1407](#) - reassign failed tablets
- [KUDU-1442](#) - log number of open log block containers
- [KUDU-1544](#) - Race in Java client's AsyncKuduSession.apply()

- [KUDU-1726](#) - Avoid fsync-per-block in tablet copy
- [KUDU-1755](#) - Part 1: Improve tablet on disk size metric
- [KUDU-1811](#) - C++ client: use larger batches when fetching scan tokens
- [KUDU-1863](#) - improve overall safety of graceful server shutdown
- [KUDU-1865](#) - Avoid heap allocation for payload slices
- [KUDU-1894](#) - fixed deadlock in client.Connection
- [KUDU-1911](#) - improve missing required arg message
- [KUDU-1929](#) - [rpc] Allow using encrypted private keys for TLS
- [KUDU-1942](#) - Kerberos fails to log in on hostnames with capital letters
- [KUDU-1943](#) - Add BlockTransaction to Block Manager
- [KUDU-1952](#) - Remove round-robin for block placement
- [KUDU-1955](#) - refuse to use world-readable keytabs
- [KUDU-2004](#) - Undefined behavior in TlsSocket::Writev()
- [KUDU-2013](#) - Support for long lived auth tokens in Java client
- [KUDU-2032](#) - Kerberos authentication fails with rdns disabled in krb5.conf
- [KUDU-2039](#) - Fix the table count in the /tables page of master webUI
- [KUDU-2041](#) - Fix negotiation deadlock
- [KUDU-2049](#) - Fix too-strict CHECK in RleIntBlockDecoder::SeekToPositionInBlock
- [KUDU-2053](#) - Fix race in Java RequestTracker
- [KUDU-2058](#) - Fix LocatedTablet string comparisons
- [KUDU-2060](#) - Show primary keys in the master's table web UI page
- [KUDU-2066](#) - Add experimental Gradle build support
- [KUDU-2067](#) - Enable cfile checksumming by default
- [KUDU-2072](#) - upgrade to cmake 3.9.0 breaks sles12sp0 cmake patch
- [KUDU-2078](#) - Sink failure if batch size > session's flush buffer size
- [KUDU-2083](#) - Decrement running maintenance ops on failed prepare
- [KUDU-2085](#) - Fix crash when seeking past end of prefix-encoded blocks
- [KUDU-2087](#) - Fix failure to map Kerberos principal to username with FreeIPA
- [KUDU-2088](#) - Synchronizer may not go out of scope with outstanding references
- [KUDU-2091](#) - Certificates with intermediate CA's do not work with Kudu
- [KUDU-2101](#) - Include a table summary at the bottom
- [KUDU-2102](#) - fix PosixRWFile::Sync to guarantee durability when used concurrently
- [KUDU-2103](#) - [java]Canonicalize hostnames in client
- [KUDU-2104](#) - Upgrade to Spark 2.2.0
- [KUDU-2114](#) - Don't re-delete tombstoned replicas
- [KUDU-2118](#) - Fully shut down TabletReplica on delete
- [KUDU-2123](#) - Auto-vivify cmeta on tombstoned replicas if doesn't exist at startup
- [KUDU-2131](#) - switch to LIFO log container retrieval
- [KUDU-2138](#) - delete failed replicas in tablet report
- [KUDU-2141](#) - master: Remove DCHECK when tablet report has no opid_index
- [OOZIE-2488](#) - Upgrade to latest OpenJPA version
- [OOZIE-2796](#) - oozie.action.keep.action.dir not getting notice
- [OOZIE-2816](#) - Strip out the first command word from Sqoop action if its "sqoop"
- [OOZIE-2854](#) - Oozie should handle transient database problems
- [OOZIE-2984](#) - Parse spark-defaults.conf values with spaces without needing the quotes
- [PARQUET-825](#) - Static analyzer findings (NPEs, resource leaks)
- [PIG-3567](#) - LogicalPlanPrinter throws OOM for large scripts
- [PIG-3655](#) - BinStorage and InterStorage approach to record markers is broken
- [SENTRY-1546](#) - Generic Policy provides bad error messages for Sentry exceptions
- [SENTRY-1580](#) - Provide pooled client connection model with HA

- [SENTRY-1646](#) - Unable to truncate table <database>.<tablename>; from "default" databases
- [SENTRY-1665](#) - cross-site scripting vulnerability in ConfServlet
- [SENTRY-1696](#) - Expose time spent creating the initial snapshot as a metric
- [SENTRY-1724](#) - Remove old PoolClientInvocationHandler
- [SENTRY-1744](#) - Simplify creation of DelegateSentryStore
- [SENTRY-1801](#) - Sentry Namenode Plugin should handle unknown permissions
- [SENTRY-1820](#) - Add JSON file reporter for Sentry metrics
- [SENTRY-1822](#) - Allow multiple Sentry reporters
- [SENTRY-1827](#) - Minimize TPathsDump thrift message used in HDFS sync
- [SENTRY-1853](#) - Add the log level access mechanism
- [SENTRY-1881](#) - PrivilegeOperatePersistence throws wrong type of exceptions
- [SENTRY-1892](#) - Reduce memory consumption of HMSPath\$Entry and TPathEntry
- [SENTRY-1903](#) - TransactionManager shows retried transactions starting from 0
- [SENTRY-1932](#) - Improve logging for HMSPath
- [SENTRY-1937](#) - Optimize performance for listing sentry roles by group name
- [SENTRY-1938](#) - Sentry logs to provide more relevant information
- [SOLR-10889](#) - Stale zookeeper information is used during failover check
- [SPARK-13278](#) - Launcher fails to start with JDK 9 EA
- [SPARK-13669](#) - Job always fails in the external shuffle service unavailable situation
- [SPARK-15067](#) - YARN executors are launched with fixed perm gen size
- [SPARK-19688](#) - Spark on YARN credentials file set to different application directory
- [SPARK-20393](#) - Strengthen Spark to prevent XSS vulnerabilities
- [SPARK-20904](#) - Task failures during shutdown cause problems with pre-empted executors
- [SQOOP-3139](#) - sqoop tries to re execute select query during import in case of a connection reset error and this is causing lots of duplicate records from source
- [ZOOKEEPER-1653](#) - zookeeper fails to start because of inconsistent epoch
- [ZOOKEEPER-2678](#) - Large databases take a long time to regain a quorum

Issues Fixed in CDH 5.12.x

The following topics describe issues fixed in CDH 5.12.x, from newest to oldest release. You can also review [What's New in CDH 5.12.x](#) on page 1819 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.12.2

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.12.2:

- [FLUME-3046](#) - Kafka Sink and Source Configuration Improvements
- [FLUME-3100](#) - Support arbitrary header substitution for topic of Kafka Sink
- [FLUME-3144](#) - Improve Log4jAppender's performance by allowing logging collection of messages
- [FLUME-3186](#) - Make asyncHbaseClient configuration parameters available from Flume configuration
- [HADOOP-13017](#) - Implementations of InputStream.read(buffer, offset, bytes) to exit 0 if bytes==0
- [HADOOP-14687](#) - AuthenticatedURL will reuse bad/expired session cookies
- [HDFS-6763](#) - Initialize file system-wide quota once on transitioning to active
- [HDFS-7582](#) - Enforce maximum number of ACL entries separately per access and default.
- [HDFS-8540](#) - Mover should exit with NO_MOVE_BLOCK if no block can be moved
- [HDFS-8541](#) - Mover should exit with NO_MOVE_PROGRESS if there is no move progress
- [HDFS-8818](#) - Allow Balancer to run faster.
- [HDFS-8865](#) - Improve quota initialization performance
- [HDFS-9003](#) - ForkJoin thread pool leaks
- [HDFS-9396](#) - Total files and directories on jmx and web UI on standby is uninitialized
- [HDFS-11384](#) - Balancer disperses getBlocks calls to avoid NameNode's rpc queue saturation

- [HDFS-11738](#) - Hedged pread takes more time when block moved from initial locations
- [HDFS-11742](#) - Improve balancer usability after HDFS-8818
- [HDFS-11755](#) - Underconstruction blocks can be considered missing
- [HDFS-11851](#) - getGlobalJNIEnv() may deadlock if exception is thrown
- [HDFS-12052](#) - Set SWEBHDFS delegation token kind when ssl is enabled in HttpFS
- [HDFS-12117](#) - HttpFS does not seem to support SNAPSHOT related methods for WebHDFS REST Interface
- [HDFS-12191](#) - Provide option to not capture the accessTime change of a file to snapshot if no other modification has been done to this file
- [HDFS-12357](#) - Let NameNode to bypass external attribute provider for configured users
- [HDFS-12363](#) - Possible NPE in BlockManager\$StorageInfoDefragmenter#scanAndCompactStorages
- [HDFS-12369](#) - Edit log corruption due to hard lease recovery of not-closed file which has snapshots.
- [HDFS-12424](#) - Datatable sorting on the Datanode Information page in the Namenode UI is broken
- [HDFS-12494](#) - libhdfs SIGSEGV in setTLSEnvironmentStrings
- [HDFS-12642](#) - Log block and datanode details in BlockRecoveryWorker.
- [YARN-6757](#) - Refactor the usage of yarn.nodemanager.linux-container-executor.cgroups.mount-path
- [YARN-7099](#) - ResourceHandlerModule.parseConfiguredCGroupPath only works for privileged yarn users.
- [YARN-7261](#) - Add debug message for better download latency monitoring.
- [HBASE-15548](#) - SyncTable: sourceHashDir is supposed to be optional but won't work without
- [HBASE-15691](#) - ConcurrentModificationException in BucketAllocator
- [HBASE-16351](#) - Improve error reporting for license failures
- [HBASE-17287](#) - Master becomes a zombie if filesystem object closes
- [HBASE-17352](#) - Fix hbase-assembly build with bash 4
- [HBASE-17534](#) - Avoid re-wrapping IOExceptions as IOExceptions
- [HBASE-17648](#) - HBase Table-level synchronization fails between two secured(kerberized) cluster
- [HBASE-17698](#) - ReplicationEndpoint choosing sinks
- [HBASE-18025](#) - CatalogJanitor should collect outdated RegionStates from the AM
- [HBASE-18030](#) - Per Cell TTL tags may get duplicated with increments/Append causing tags length overflow.
- [HBASE-18036](#) - Data locality is not maintained after cluster restart or SSH
- [HBASE-18066](#) - Get with closest_row_before on hbase:meta can return empty Cell during region merge/split
- [HBASE-18077](#) - Update JUnit licensing to use EPL
- [HBASE-18081](#) - The way we process connection preamble in SimpleRpcServer is broken
- [HBASE-18122](#) - Scanner id should include ServerName of region server
- [HBASE-18390](#) - Sleep too long when finding region location failed
- [HBASE-18487](#) - Minor fixes in row lock implementation
- [HBASE-18577](#) - Shaded client includes several non-relocated third party dependencies
- [HBASE-18665](#) - ReversedScannerCallable invokes getRegionLocations incorrectly
- [HBASE-18675](#) - Making {max,min}SessionTimeout configurable for MiniZooKeeperCluster
- [HBASE-18924](#) - Backport HBASE-18568 (Correct metric of numRegions) to branch-1.2 and branch-1.3
- [HBASE-19215](#) - Incorrect exception handling on the client causes incorrect call timeouts and byte buffer allocations on the server
- [HIVE-6727](#) - Table level stats for external tables are set incorrectly
- [HIVE-9447](#) - Metastore: inefficient Oracle query for removing unused column descriptors when add/drop table/partition
- [HIVE-15233](#) - UDF UUID() should be non-deterministic
- [HIVE-15519](#) - BitSet not computed properly for ColumnBuffer subset
- [HIVE-16758](#) - Better Select Number of Replications
- [HIVE-16784](#) - Missing lineage information when hive.blobstore.optimizations.enabled is true
- [HIVE-16949](#) - Leak of threads from Get-Input-Paths and Get-Input-Summary thread pool
- [HIVE-17022](#) - Add mode in lock debug statements
- [HIVE-17218](#) - Canonical-ize hostnames for Hive metastore, and HS2 servers.

- [HIVE-17274](#) - RowContainer spills for timestamp column throws exception
- [HIVE-17640](#) - Comparison of date return null if time part is provided in string.
- [HIVE-17805](#) - SchemaTool validate locations should not return exit 1
- [HIVE-17969](#) - Metastore to alter table in batches of partitions when renaming table
- [HIVE-18009](#) - Multiple lateral view query is slow on hive on spark
- [HUE-6045](#) - [Pig] Update LOG_END_PATTERN to identify completed jobs
- [HUE-6851](#) - [metastore] Don't show the database while loading a table
- [HUE-6901](#) - [oozie] Editor is missing regular ssh action
- [HUE-7106](#) - [oozie] New oozie doc url doesn't update after saving the page
- [HUE-7127](#) - [backend] Fix for Python requests Kerberos/GSSAPI authentication library fails to authenticate kerberos requests to the same destination host.
- [HUE-7141](#) - [dashboard] Menu shows an incomplete dropdown
- [HUE-7155](#) - [backend] Hue needs to support TLSv1_1 and TLSv1_2 via Thrift connections
- [HUE-7179](#) - [core] Avoid issue with searching documents on assist or global search box
- [HUE-7182](#) - [editor] Accented characters not inserted correctly
- [HUE-7192](#) - [desktop] Hue Search is failing with DatabaseError field_cast_sql deprecated after Django 1.8
- [HUE-7223](#) - [autocomplete] Fix issue where the old autocompleter inserts extra characters on exception
- [HUE-7225](#) - [autocomplete] Prevent new line on enter when a suggestion equals the typed text
- [HUE-7264](#) - [core] Fix cx_Oracle removed numbersAsStrings in 6.0
- [HUE-7283](#) - [editor] Fix issue where analytic functions aren't found in the context popover
- [HUE-7335](#) - [editor] Make it possible to close the editor sessions panel form within the panel
- [HUE-7390](#) - [autocomplete] Fall back to only autocomplete the active statement in case there are surrounding errors
- [HUE-7402](#) - [core] Truncate SEARCH Column in Document2 for Oracle database
- [HUE-7412](#) - [editor] Fix indexes for chart limit array slice
- [HUE-7446](#) - [core] RM HA gets called by mistake on each navopt call
- [HUE-7456](#) - [core] Do not retry Thrift call on read operation timed out from ssl.SSLError
- [HUE-7470](#) - [editor] Only check for risks when a snippet has focus
- [HUE-7476](#) - [core] Add Hue shell command to set CM environment
- [HUE-7489](#) - [editor] CloseOperation when GetOperationStatus is in an ERROR_STATE
- [HUE-7519](#) - [assist] Fix JS error when the sample tab can't load samples
- [HUE-7521](#) - [security] Prevent JS error on printAcl not defined
- [HUE-7525](#) - [home] Prevent js error when dragging to select documents
- [HUE-7526](#) - [frontend] Avoid XSS in the wysiwyg5 editor and fix the binding in Hue 4
- [HUE-7530](#) - [editor] Avoid red error popup on close of expired query
- [HUE-7533](#) - [fb] Wrong error message in UI with 500 error when user tries to edit read_only file in File browser
- [HUE-7536](#) - [frontend] Avoid XSS in the table row display modal, HueDatatable and TableExtender2
- [HUE-7538](#) - [frontend] Override default knockout html binding to prevent XSS
- [HUE-7552](#) - [fb] Copy Folder Into Itself Causes Endless Loop
- [HUE-7564](#) - Remove one of the get_connection(ldap_config) function in ldap_access.py.
- [HUE-7567](#) - [all] Add default timeout to REST calls
- [HUE-7592](#) - [jib] Hide older browser if the new ones is enabled
- [HUE-7622](#) - [jib] Reruning a wokflow redirects to a json result
- [IMPALA-4682](#) - Remove Preconditions check from analyzeAggregation().
- [IMPALA-4737](#) - Prevent SIGUSR1 from killing daemons when minidumps are disabled
- [IMPALA-4826](#) - Fix error during a scan on repeated root schema in Parquet.
- [IMPALA-4951](#) - Fix database visibility for user with only column privilege
- [IMPALA-5327](#) - Handle return of JNI GetStringUTFChar
- [IMPALA-5336](#) - Fix partition pruning when column is cast
- [IMPALA-5446](#) - dropped Sorter::Reset() status

- [IMPALA-5495](#) - Improve error message if no impalad role is configured
- [IMPALA-5504](#) - Fix TupleIsNullPredicate evaluation.
- [IMPALA-5531](#) - Fix correctness issue in correlated aggregate subqueries
- [IMPALA-5595](#) - Only set KuduScanner timestamp feature flag if necessary
- [IMPALA-5597](#) - Try casting targetExpr when building runtime filter plan
- [IMPALA-5598](#) - Fix excessive dumping in MemLimitExceeded
- [IMPALA-5689](#) - Avoid inverting non-equi left joins
- [IMPALA-5750](#) - Catch exceptions from boost thread creation
- [IMPALA-5776](#) - Write partial tuple to the correct mempool
- [IMPALA-5784](#) - Separate planner and user set query options in profile
- [IMPALA-5796](#) - CTAS for Kudu fails with expr rewrite
- [IMPALA-5815](#) - right outer join returns invalid memory
- [IMPALA-5819](#) - DCHECK in HdfsTextScanner::Close()
- [IMPALA-5820](#) - Fix string format() syntax in test_scanners_fuzz.py
- [IMPALA-5840](#) - Don't write page-level statistics in Parquet files.
- [IMPALA-5850](#) - Cast sender partition exprs under unions.
- [IMPALA-5856](#) - Fix outer join predicate assignment.
- [IMPALA-5871](#) - KuduPartitionExpr incorrectly handles its child types
- [IMPALA-5885](#) - free runtime filter allocations in Parquet
- [IMPALA-5892](#) - Allow reporting status independent of fragment instance
- [IMPALA-5923](#) - Print binary ID as hex in ChildQuery::Cancel()
- [IMPALA-5954](#) - Set DO_NOT_UPDATE_STATS in alterTable() RPCs to HMS.
- [IMPALA-5983](#) - Fix crash in to/from_utc_timestamp
- [IMPALA-5994](#) - Lower case struct-field names
- [IMPALA-6060](#) - Check the return value of JNI exception handling functions
- [KITE-762](#) - Multiple URIs in hive.metastore.uris configuration may be problematic for Crunch+Kite
- [OOZIE-1401](#) - PurgeCommand should purge the workflow jobs w/o end_time
- [OOZIE-3103](#) - Improve log message regarding max concurrency for coordinators
- [SENTRY-1768](#) - Avoid detaching object on transaction exit when it isn't needed
- [SENTRY-1770](#) - Avoid more detaches on commit
- [SENTRY-1892](#) - Reduce memory consumption of HMSPath\$Entry and TPathEntry
- [SENTRY-1966](#) - Improve logging of HMS sync data (paths and permissions) flowing from Sentry to NameNode
- [SENTRY-1993](#) - StringIndexOutOfBoundsException in HMSPathsDumper.java
- [SENTRY-2027](#) - backportCreate mechanism of delivering commands via WebUI
- [SENTRY-2066](#) - DB name is not set for AlterTable
- [SPARK-17460](#) - [SQL] Make sure sizeInBytes in Statistics will not overflow
- [SPARK-17511](#) - Yarn Dynamic Allocation: Avoid marking released container as Failed
- [SPARK-17531](#) - Fix the "hive.metastore.uris" value that was changed by the backport of to 1.6 branches.
- [SQOOP-3014](#) - Sqoop with HCatalog import loose precision for large numbers that does not fit into double
- [SQOOP-3139](#) - sqoop tries to re execute select query during import in case of a connection reset error and this is causing lots of duplicate records from source
- [SQOOP-3243](#) - Importing BLOB data causes 'Stream closed' error on encrypted HDFS
- [ZOOKEEPER-2347](#) - Deadlock shutting down zookeeper
- [ZOOKEEPER-2678](#) - Discovery and Sync can take a very long time on large DB
- [ZOOKEEPER-2726](#) - backportPatch for ZOOKEEPER-2693 introduces potential race condition

Issues Fixed in CDH 5.12.1

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.12.1:

- [FLUME-2752](#) - Fix AvroSource startup resource leaks

- [FLUME-2905](#) - Fix NetcatSource file descriptor leak if startup fails
- [HADOOP-13588](#) - ConfServlet should respect Accept request header
- [HADOOP-13628](#) - Support to retrieve specific property from configuration via REST API
- [HADOOP-14260](#) - Configuration. dumpConfiguration should redact sensitive information
- [HADOOP-14511](#) - WritableRpcEngine.Invocation#toString NPE on null parameters
- [HADOOP-14542](#) - Add IOUtils. cleanupWithLogger that accepts slf4j logger API
- [HDFS-8856](#) - Make LeaseManager#countPath O(1).
- [HDFS-10468](#) - HDFS read ends up ignoring an interrupt
- [HDFS-10506](#) - OIV's ReverseXML processor cannot reconstruct some snapshot details
- [HDFS-11303](#) - Hedged read might hang infinitely if read data from all DN failed
- [HDFS-11708](#) - Positional read will fail if replicas moved to different DNs after stream is opened
- [HDFS-11741](#) - Long running balancer may fail due to expired DataEncryptionKey
- [HDFS-11861](#) - ipc.Client. Connection#sendRpcRequest should log request name
- [HDFS-11881](#) - NameNode consumes a lot of memory for snapshot diff report generation
- [HDFS-11960](#) - Successfully closed files can stay under-replicated
- [HDFS-12042](#) - Lazy initialize AbstractINodeDiffList#diffs for snapshots to reduce memory consumption
- [HDFS-12139](#) - HTTPFS liststatus returns incorrect pathSuffix for path of file
- [HDFS-12278](#) - LeaseManager operations are inefficient in 2.8
- [MAPREDUCE-6870](#) - Add configuration for MR job to finish when all reducers are complete.
- [YARN-2780](#) - Log aggregated resource allocation in rm-appsummary.log
- [HBASE-15720](#) - Print row locks at the debug dump page
- [HBASE-16033](#) - Add more details in logging of responseTooSlow/TooLarge
- [HBASE-17131](#) - Avoid livelock caused by HRegion#processRowsWithLocks
- [HBASE-17587](#) - Do not Rethrow DoNotRetryIOException as UnknownScannerException
- [HBASE-18247](#) - Hbck to fix the case that replica region shows as key in the meta table
- [HBASE-18362](#) - hbck should not report split replica parent region from meta as errors
- [HIVE-9567](#) - JSON SerDe not escaping special chars when writing char/varchar data
- [HIVE-10209](#) - FetchTask with VC may fail because ExecMapper.done is true
- [HIVE-11240](#) - Change value type from int to long for HiveConf.ConfVars.METASTORESERVERMAXMESSAGEIZE
- [HIVE-11462](#) - GenericUDFStruct should constant fold at compile time.
- [HIVE-11592](#) - ORC metadata section can sometimes exceed protobuf message size limit
- [HIVE-12274](#) - Increase width of columns used for general configuration in the metastore.
- [HIVE-12551](#) - Fix several kryo exceptions in branch-1
- [HIVE-12762](#) - Common join on parquet tables returns incorrect result when hive.optimize.index. filter set to true
- [HIVE-13330](#) - ORC vectorized string dictionary reader does not differentiate null vs empty string dictionary
- [HIVE-13588](#) - NPE is thrown from MapredLocalTask.executeInChildVM
- [HIVE-13947](#) - HoS prints wrong number for hash table size in map join scenario
- [HIVE-14178](#) - Hive::needsToCopy should reuse FileUtils::equalsFileSystem
- [HIVE-15122](#) - Hive: Upcasting types should not obscure stats (min/max/ndv)
- [HIVE-15792](#) - Hive should raise SemanticException when LPAD/RPAD pad character's length is 0
- [HIVE-16183](#) - Fix potential thread safety issues with static variables
- [HIVE-16291](#) - Hive fails when unions a parquet table with itself
- [HIVE-16559](#) - Parquet schema evolution for partitioned tables may break if table and partition serdes differ
- [HIVE-16729](#) - Improve location validator to check for blank paths
- [HIVE-16845](#) - INSERT OVERWRITE a table with dynamic partitions on S3 fails with NPE
- [HIVE-16869](#) - Hive returns wrong result when predicates on non-existing columns are pushed down to Parquet reader
- [HIVE-16875](#) - Query against view with partitioned child on HoS fails with privilege exception.
- [HIVE-16930](#) - HoS should verify the value of Kerberos principal and keytab file before adding them to spark-submit command parameters

- [HIVE-16935](#) - Hive should strip comments from input before choosing which CommandProcessor to run.
- [HIVE-16974](#) - Change the sort key for the schema tool validator to be <ID>
- [HIVE-17050](#) - Multiline queries that have comment in middle fail when executed via "beeline -e"
- [HIVE-17052](#) - Remove logging of predicate filters
- [HIVE-17149](#) - Hdfs directory is not cleared if partition creation failed on HMS
- [HUE-5217](#) - [backend] Resolve LGPL copyleft issue - Paramiko in boto 2.46.1
- [HUE-5504](#) - [oozie] Only use JDBC URL from hive2 action when hardcoded
- [HUE-6684](#) - [editor] The horizontal scrollbar of the ace editor isn't updated when the right assist is resized.
- [HUE-6701](#) - [editor] The right panel resize bar gets hidden in Hue 3 on window resize.
- [HUE-6702](#) - [about] About page is accessible without authentication
- [HUE-6704](#) - [fb] Handle empty page error on filebrowser filter requests
- [HUE-6705](#) - [fb] Opening directory after opening editor does not load at first click
- [HUE-6709](#) - [editor] Column list can jump back and forth depending on its size
- [HUE-6710](#) - [notebook] Application reachable directly by users without granted access
- [HUE-6711](#) - [frontend] Render 403 page properly on Hue 4
- [HUE-6714](#) - [job] No tasks found for job when refreshing MapReduce job page
- [HUE-6724](#) - [job] Job progress can have too many decimals
- [HUE-6730](#) - [fb] Use SkipFile instead of StopUpload in case of exception during file upload
- [HUE-6735](#) - [job] A suspended workflow should also be killable
- [HUE-6738](#) - [job] Kill button of workflow enabled when clicking on re-run
- [HUE-6742](#) - [editor] Incorrect Sorting With TIMESTAMP Column Type
- [HUE-6745](#) - [frontend] The title of the browser is Filebrowser even if I'm in editor
- [HUE-6748](#) - [editor] Select and copy the query from the query history
- [HUE-6749](#) - [editor] From metastore to editor with past query ran and scrolled can garble grid
- [HUE-6758](#) - [job] Open a workflow workspace breaks the history
- [HUE-6760](#) - [job] Filtering running jobs and batch killing them do not update their status
- [HUE-6767](#) - [job] Submit a workflow sometimes generates a JS error
- [HUE-6768](#) - [frontend] Remove shown tooltips on app change in Hue 4
- [HUE-6772](#) - [metadata] Cache Navigator password reads from scripts
- [HUE-6776](#) - [search] Js error when clicking on the first two layout of a new dashboard
- [HUE-6777](#) - [fb] Drag & drop upload not enabled if editor was loaded before
- [HUE-6779](#) - [job] Disable click around checkbox in list of jobs
- [HUE-6780](#) - [s3] Correctly infer and display region when connected to S3 by endpoint
- [HUE-6781](#) - [job] Workflow duration time is not humanized
- [HUE-6786](#) - [metastore] Fix ko editable binding after applying toggle overflow to the db description
- [HUE-6789](#) - [beeswax] Fix alignment of the assist headers
- [HUE-6791](#) - [search] Protect against pivot facets conflicting with nested facets
- [HUE-6792](#) - [core] Middle/right click is not working as a link on main blue action button
- [HUE-6795](#) - [fb] Fix minor logic error in trash directory creation
- [HUE-6797](#) - [dashboard] Grid widget operation icons are not grayed much
- [HUE-6800](#) - [frontend] Re-word the welcome tour
- [HUE-6802](#) - [editor] Don't check for column existence in the location handler when a column is prefixed with a table name or alias
- [HUE-6810](#) - [assist] Improve positioning and visibility of side panel toggles
- [HUE-6813](#) - [metastore] Fix database comment last character truncation
- [HUE-6814](#) - [search] Only return distinct usernames in security impersonate dropdown
- [HUE-6819](#) - [oozie] Set generic widget for generic actions in oozie graph
- [HUE-6820](#) - [frontend] Catch all the JS errors and send them to the backend
- [HUE-6825](#) - [autocomplete] Prevent exception when there's no identifierChain given for a table
- [HUE-6833](#) - [editor] Disable query builder tab for non-sql editors

- [HUE-6840](#) - [notebook] Workaround lack of status flag in Impala profile page
- [HUE-6842](#) - [core] Correctly append the list of custom apps
- [HUE-6843](#) - [metastore] Show partition icon on table browser sends to Hue 3
- [HUE-6844](#) - [frontend] Make all custom apps embeddable
- [HUE-6845](#) - [frontend] Fix custom app mako generation
- [HUE-6846](#) - [frontend] Make all proxy apps embeddable
- [HUE-6847](#) - [metastore] Do not display if table is compressed as Hive information is incorrect
- [HUE-6856](#) - [search] Protect against reflected XSS in search query parameters
- [HUE-6863](#) - [job] Kill buttons are enabled even when the user is not allowed
- [HUE-6864](#) - [job] Check for permission before killing a job
- [HUE-6870](#) - [job] Improve end user UX with bubbling up of common actions and properties.
- [HUE-6883](#) - [metastore] Allow deletion of table comments
- [HUE-6885](#) - [frontend] Catch all http 502 and show an HTML stripped version of the message
- [HUE-6886](#) - [oozie] SLA variable in workflow action are not getting retrieved
- [HUE-6888](#) - [oozie] Fix link page on SLA with Hue 4
- [HUE-6894](#) - [core] Remove cluster config call on login page
- [HUE-6896](#) - [frontend] Errorcatcher should not log jQuery triggered events
- [HUE-6897](#) - [oozie] Ensure that examples_dir exists before installing examples
- [HUE-6898](#) - [oozie] Editor doesn't update job status to "finished" for finished Pig jobs
- [HUE-6899](#) - [core] Allow 403 to be accessed by regular users
- [HUE-6900](#) - [useradmin] Do not link app title to list of users if user not a super user
- [HUE-6901](#) - [oozie] Editor is missing certain regular actions
- [HUE-6903](#) - [dataeng] Fix for s3 context popover
- [HUE-6905](#) - [frontend] Minor UI bugs in old Pig Editor and Job Designer.
- [HUE-6939](#) - [editor] Improve interactions with the risk indicator
- [HUE-6940](#) - [editor] Cancel risk requests if editing while risk check is running
- [HUE-6941](#) - [editor] Don't show a pointer cursor for the risk check spinner
- [HUE-6950](#) - [sam] Create home directory for new user login
- [HUE-6955](#) - [metadata] Extend query suggest caching to 1 week
- [HUE-6960](#) - [frontend] Only suggest facet values and no results when autocompleting facets in the top search
- [HUE-6995](#) - [oozie] Set minimum width of a workflow node
- [HUE-7000](#) - [assist] HBase assist does not load the table the second time
- [HUE-7001](#) - [assist] Remember the last selected HBase cluster
- [HUE-7002](#) - [core] Welcome to Hue 4 tour could be also skipped by clicking on the white background
- [HUE-7004](#) - [notebook] Loading query with empty session param shows several JS errors
- [HUE-7005](#) - [oozie] sharing coordinator or bundle immediately after creating fails
- [HUE-7006](#) - [core] Ability in Hue to change Cookie name from default "sessionid"
- [HUE-7040](#) - [oozie] Pressing Esc on share popup throws JS error in Hue4
- [HUE-7045](#) - [hive] Avoid use database call for SET queries
- [HUE-7048](#) - [oozie] User prompting doesn't work in share document popup in hue4
- [HUE-7052](#) - [desktop] Retain last_modified when sharing sample docs in sync_documents
- [HUE-7069](#) - [editor] Remove the spinner for risk checking
- [HUE-7080](#) - [oozie] Oozie HiveServer2 action is prompting for and 'Impalad hostname'
- [HUE-7082](#) - [aws] Create S3 bucket fails with error Bad request
- [HUE-7084](#) - [useradmin] Remove @domain part of username when creating home directory for LDAP user
- [HUE-7090](#) - [security] Protect from missing groupName in the list_sentry_roles_by_group API
- [HUE-7104](#) - [core] Hue and Isilon issue - UnsupportedOperationException: Unknown op 'GETTRASHROOT'
- [HUE-7107](#) - [core] Hue 4 group access doesn't reflect on Interface
- [IMPALA-4276](#) - Profile displays non-default query options set by planner
- [IMPALA-4866](#) - Hash join node does not apply limits correctly

- [IMPALA-5354](#) - INSERT hints for Kudu tables
- [IMPALA-5427](#) - Fix race between CRS::UpdateQueryStatus() and beeswax RPCs
- [IMPALA-5500](#) - Reduce catalog update topic size
- [IMPALA-5524](#) - Fixes NPE during planning with DISABLE_UNSAFE_SPILLS=1
- [IMPALA-5539](#) - Fix Kudu timestamp with -use_local_tz_for_unix_ts
- [IMPALA-5554](#) - sorter DCHECK on null column
- [IMPALA-5567](#) - race in fragment instance teardown
- [IMPALA-5579](#) - Fix IndexOutOfBoundsException in GetTables metadata request
- [IMPALA-5580](#) - fix Java UDFs that return NULL strings
- [IMPALA-5582](#) - Store sentry privileges in lower case
- [IMPALA-5586](#) - Null-aware anti-join can take a long time to cancel
- [IMPALA-5588](#) - Reduce the frequency of fault injection
- [IMPALA-5611](#) - KuduPartitionExpr holds onto memory unnecessarily
- [IMPALA-5615](#) - Fix compute incremental stats for general partition exprs
- [IMPALA-5616](#) - Add --enable_minidumps startup flag
- [IMPALA-5623](#) - Fix lag() on STRING cols to release UDF mem
- [IMPALA-5627](#) - fix dropped statuses in HDFS writers
- [IMPALA-5638](#) - Fix Kudu table set tblproperties inconsistencies
- [IMPALA-5657](#) - Fix a couple of bugs with FunctionCallExpr and IGNORE NULLS
- [IMPALA-5686](#) - Update a mini cluster Sentry property
- [IMPALA-5691](#) - recalibrate mem limit for Q18
- [KITE-1155](#) - Deleting an already deleted empty path should not fail the job
- [OOZIE-2816](#) - Strip out the first command word from Sqoop action if its "sqoop"
- [OOZIE-2984](#) - Parse spark-defaults.conf values with spaces without needing the quotes
- [PIG-3567](#) - LogicalPlanPrinter throws OOM for large scripts
- [PIG-3655](#) - BinStorage and InterStorage approach to record markers is broken
- [SENTRY-1646](#) - Unable to truncate table <database>.<tablename>; from "default" databases
- [SENTRY-1811](#) - Optimize data structures used in HDFS sync
- [SENTRY-1827](#) - Minimize TPathsDump thrift message used in HDFS sync
- [SOLR-6673](#) - MDC based logging of collection, shard, etc.
- [SOLR-10889](#) - Stale zookeeper information is used during failover check.
- [SPARK-13278](#) - [CORE] Launcher fails to start with JDK 9 EA
- [SPARK-13669](#) - [SPARK-20898] [CORE] Improve the blacklist mechanism to handle external shuffle service unavailable situation
- [SPARK-15067](#) - [YARN] YARN executors are launched with fixed perm gen size
- [SPARK-16845](#) - [SQL][BRANCH-1.6] GeneratedClass\$SpecificOrdering` grows beyond 64 KB
- [SPARK-19019](#) - [PYTHON] Fix hijacked `collections.namedtuple` and port cloudpickle changes for PySpark to work with Python 3.6.0
- [SPARK-19688](#) - [STREAMING] Not to read `spark.yarn.credentials.file` from checkpoint.
- [SPARK-20393](#) - [WEBUI] Strengthen Spark to prevent XSS vulnerabilities
- [SPARK-20904](#) - [CORE] Don't report task failures to driver during shutdown.
- [ZOOKEEPER-1653](#) - zookeeper fails to start because of inconsistent epoch
- [ZOOKEEPER-2040](#) - Server to log underlying cause of SASL connection problems.

Issues Fixed in CDH 5.12.0

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.12.0:

- [CRUNCH-636](#) - Make replication factor for temporary files configurable
- [FLUME-3092](#) - Extend the FileChannel's monitoring metrics
- [HADOOP-11379](#) - Fix new findbugs warnings in hadoop-auth*

- [HADOOP-11385](#) - Prevent cross site scripting attack on JMXJSONServlet
- [HADOOP-11628](#) - SPNEGO auth does not work with CNAMES in JDK8.
- [HADOOP-11709](#) - Time.NANOSECONDS_PER_MILLISECOND - use class-level final constant instead of method variable
- [HADOOP-11794](#) - Enable distcp to copy blocks in parallel
- [HADOOP-11812](#) - Implement listLocatedStatus for ViewFileSystem to speed up split calculation
- [HADOOP-11962](#) - Sasl message with MD5 challenge text shouldn't be LOG out even in debug level
- [HADOOP-12054](#) - RPC client should not retry for InvalidToken exceptions.
- [HADOOP-12469](#) - distcp should not ignore the ignoreFailures option
- [HADOOP-12622](#) - Improve the loggings in RetryPolicies and RetryInvocationHandler
- [HADOOP-12751](#) - While using kerberos Hadoop incorrectly assumes names with '@' to be non-simple.
- [HADOOP-12906](#) - AuthenticatedURL should convert a 404/Not Found into an FileNotFoundException.
- [HADOOP-13138](#) - Unable to append to a SequenceFile with Compression.NONE
- [HADOOP-13163](#) - Reuse pre-computed filestatus in Distcp-CopyMapper
- [HADOOP-13169](#) - Randomize file list in SimpleCopyListing
- [HADOOP-13174](#) - Add more debug logs for delegation tokens and authentication.
- [HADOOP-13441](#) - Document LdapGroupsMapping keystore password properties
- [HADOOP-13453](#) - S3Guard: Instrument new functionality with Hadoop metrics
- [HADOOP-13552](#) - RetryInvocationHandler logs all remote exceptions
- [HADOOP-13626](#) - Remove distcp dependency on FileStatus serialization
- [HADOOP-13720](#) - Add more info to the msgs printed in AbstractDelegationTokenSecretManager
- [HADOOP-13760](#) - S3Guard: add delete tracking
- [HADOOP-13854](#) - KMS should log error details in KMSExceptionsProvider.
- [HADOOP-14036](#) - S3Guard: intermittent duplicate item keys failure
- [HADOOP-14081](#) - S3A: Consider avoiding array copy in S3ABlockOutputStream (ByteArrayBlock)
- [HADOOP-14100](#) - Upgrade Jsch jar to latest version to fix vulnerability in old versions
- [HADOOP-14120](#) - needless S3AFileSystem.setOptionalPutRequestParameters in S3ABlockOutputStream putObject()
- [HADOOP-14123](#) - Remove misplaced ADL service provider config file for FileSystem
- [HADOOP-14135](#) - Remove URI parameter in AWSCredentialProvider constructors
- [HADOOP-14138](#) - Remove S3A ref from META-INF service discovery, rely on existing core-default entry
- [HADOOP-14141](#) - Store KMS SSL keystore password in catalina.properties
- [HADOOP-14153](#) - ADL module has messed doc structure
- [HADOOP-14173](#) - Remove unused AdlConfKeys#ADL_EVENTS_TRACKING_SOURCE
- [HADOOP-14196](#) - Azure Data Lake doc is missing required config entry
- [HADOOP-14197](#) - Fix ADLS doc for credential provider
- [HADOOP-14214](#) - DomainSocketWatcher::add()/delete() should not self interrupt while looping await()
- [HADOOP-14241](#) - Add ADLS sensitive config keys to default list
- [HADOOP-14242](#) - Make KMS Tomcat SSL property sslEnabledProtocols and clientAuth configurable
- [HADOOP-14246](#) - Authentication Tokens should use SecureRandom instead of Random and 256 bit secrets
- [HADOOP-14266](#) - S3Guard: S3AFileSystem::listFiles() to employ MetadataStore
- [HADOOP-14369](#) - NetworkTopology calls expensive toString() when logging
- [HDFS-2219](#) - Change fsck to support fully qualified paths so that a particular namenode in a federated cluster with multiple namenodes can be specified in the path parameter.
- [HDFS-6757](#) - Simplify lease manager with INodeID
- [HDFS-7433](#) - Optimize performance of DatanodeManager's node map
- [HDFS-7939](#) - Two fsimage_rollback_* files are created which are not deleted after rollback.
- [HDFS-7959](#) - WebHdfs logging is missing on Datanode
- [HDFS-8498](#) - Blocks can be committed with wrong size
- [HDFS-9670](#) - DistCp throws NPE when source is root.
- [HDFS-10220](#) - A large number of expired leases can make namenode unresponsive and cause failover

- [HDFS-10301](#) - Remove FBR tracking state to fix false zombie storage detection for interleaving block reports
- [HDFS-10987](#) - Make Decommission less expensive when lot of blocks present
- [HDFS-11198](#) - NN UI should link DN web address using hostnames
- [HDFS-11340](#) - DataNode reconfigure for disks doesn't remove the failed volumes.
- [HDFS-11373](#) - Backport HDFS-11258 and HDFS-11272(File mtime change could not save to editlog).
- [HDFS-11383](#) - Intern strings in BlockLocation and ExtendedBlock
- [HDFS-11391](#) - Numeric usernames do no work with WebHDFS FS write access.
- [HDFS-11402](#) - HDFS Snapshots should capture point-in-time copies of OPEN files.
- [HDFS-11421](#) - Make WebHDFS' ACLs RegEx configurable
- [HDFS-11529](#) - Add libHDFS API to return last exception
- [HDFS-11531](#) - Expose hedged read metrics via libHDFS API
- [HDFS-11579](#) - Make HttpFS Tomcat SSL property sslEnabledProtocols and clientAuth configurable
- [HDFS-11609](#) - Some blocks can be permanently lost if nodes are decommissioned while dead
- [HDFS-11687](#) - Add new public encryption APIs required by Hive.
- [HDFS-11702](#) - Remove indefinite caching of key provider uri in DFSClient
- [HDFS-11709](#) - StandbyCheckpoint should handle a non-existing legacyOivlImageDir gracefully
- [HDFS-11710](#) - hadoop-hdfs-native-client build fails in trunk in Windows after HDFS-11529
- [HDFS-11714](#) - Newly added NN storage directory won't get initialized and cause space exhaustion
- [HDFS-11724](#) - libhdfs compilation is broken on OS X
- [HDFS-11914](#) - Add more diagnosis info for fsimage transfer failure
- [MAPREDUCE-5485](#) - Allow repeating job commit by extending OutputCommitter API
- [MAPREDUCE-5653](#) - DistCp does not honour config-overrides for mapreduce.[map,reduce].memory.mb
- [MAPREDUCE-6478](#) - Add an option to skip cleanupJob stage or ignore cleanup failure during commitJob.
- [MAPREDUCE-6595](#) - Fix findbugs warnings in OutputCommitter and FileOutputCommitter
- [MAPREDUCE-6762](#) - ControlledJob#toString failed with NPE when job status is not successfully updated
- [MAPREDUCE-6825](#) - YARNRunner#createApplicationSubmissionContext method is longer than 150 lines
- [MAPREDUCE-6850](#) - Shuffle Handler keep-alive connections are closed from the server side
- [MAPREDUCE-6852](#) - Job#updateStatus() failed with NPE due to race condition
- [MAPREDUCE-6871](#) - Allow users to specify racks and nodes for strict locality for AMs
- [YARN-3078](#) - LogCLIHelpers lacks of a blank space before string 'does not exist'
- [YARN-3742](#) - YARN RM will shut down if ZKClient creation times out.
- [YARN-4593](#) - Deadlock in AbstractService.getConfig()
- [YARN-6050](#) - AMs can't be scheduled on racks or nodes
- [YARN-6368](#) - Decommissioning an NM results in a -1 exit code
- [YARN-6380](#) - FSAppAttempt keeps redundant copy of the queue
- [YARN-6381](#) - FSAppAttempt has several variables that should be final
- [YARN-6510](#) - Fix profs stat file warning caused by process names that includes parenthesis.
- [YARN-6615](#) - AmpFilter drops query parameters on redirect
- [HBASE-14983](#) - Create metrics for per block type hit/miss ratios
- [HBASE-15353](#) - Add metric for number of CallQueueTooBigException's
- [HBASE-15376](#) - ScanNext metric is size-based while every other per-operation metric is time based
- [HBASE-15377](#) - Per-RS Get metric is time based, per-region metric is size-based
- [HBASE-15464](#) - Flush / Compaction metrics revisited
- [HBASE-15780](#) - Make AuthUtil public and rely on it for talking to secure HBase.
- [HBASE-15837](#) - Memstore size accounting is wrong if postBatchMutate() throws exception
- [HBASE-15941](#) - HBase repair should not unsplit healthy splitted region
- [HBASE-16052](#) - Improve HBaseFsck Scalability
- [HBASE-16350](#) - Undo server abort from HBASE-14968
- [HBASE-16561](#) - Add metrics about read/write/scan queue length and active read/write/scan handler count
- [HBASE-16630](#) - Fragmentation in long running Bucket Cache

- [HBASE-16739](#) - Timed out exception message should include encoded region name
- [HBASE-16977](#) - VerifyReplication should log a printable representation of the row keys
- [HBASE-17460](#) - enable_table_replication can not perform cyclic replication of a table
- [HBASE-17501](#) - guard against NPE while reading FileTrailer and HFileBlock
- [HBASE-17673](#) - Monitored RPC Handler not shown in the WebUI
- [HBASE-17682](#) - Region stuck in merging_new state indefinitely
- [HBASE-17688](#) - MultiRowRangeFilter not working correctly if given same start and stop RowKey
- [HBASE-17710](#) - HBase in standalone mode creates directories with 777 permission
- [HBASE-17712](#) - Remove/Simplify the logic of RegionScannerImpl.handleFileNotFound
- [HBASE-17717](#) - Incorrect ZK ACL set for HBase superuser
- [HBASE-17722](#) - Metrics subsystem stop/start messages add a lot of useless bulk to operational logging
- [HBASE-17729](#) - Missing shortcuts for some useful HCD options
- [HBASE-17731](#) - Fractional latency reporting in MultiThreadedAction
- [HBASE-17736](#) - Some options can't be configured by the shell
- [HBASE-17779](#) - disable_table_replication returns misleading message and does not turn off replication
- [HBASE-17780](#) - BoundedByteBufferPool "At capacity" messages are not actionable
- [HBASE-17798](#) - RpcServer.Listener.Reader can abort due to CancelledKeyException
- [HBASE-17803](#) - PE always re-creates table when we specify the split policy
- [HBASE-17877](#) - Improve HBase's byte[] comparator.
- [HBASE-17970](#) - Set yarn.app.mapreduce.am.staging-dir when starting MiniMRCluster
- [HBASE-18096](#) - Limit HFileUtil visibility and add missing annotations
- [HIVE-1575](#) - get_json_object does not support JSON array at the root level
- [HIVE-9679](#) - Remove redundant null-checks from DbNotificationListener
- [HIVE-9954](#) - UDFJson uses the == operator to compare Strings
- [HIVE-11418](#) - Dropping a database in an encryption zone with CASCADE and trash enabled fails
- [HIVE-11487](#) - Add getNumPartitionsByFilter api in metastore api
- [HIVE-11878](#) - ClassNotFoundException can possibly occur if multiple jars are registered one at a time in Hive
- [HIVE-12065](#) - FS stats collection may generate incorrect stats for multi-insert query
- [HIVE-13183](#) - Adds MapredLocalTasks logs in operation log
- [HIVE-13780](#) - Allow user to update AVRO table schema via command even if table's definition was defined through schema file
- [HIVE-13782](#) - Compile async query asynchronously
- [HIVE-14080](#) - hive.metastore.schema.verification should check for schema compatibility
- [HIVE-14340](#) - Add a new hook triggers before query compilation and after query execution
- [HIVE-14559](#) - Remove setting hive.execution.engine in qfiles
- [HIVE-14564](#) - Column Pruning generates out of order columns in SelectOperator which cause ArrayIndexOutOfBoundsException.
- [HIVE-14774](#) - Canceling query using Ctrl-C in beeline might lead to stale locks
- [HIVE-14943](#) - Partial Backport of HIVE-14943 - Base Implementation (of HIVE-10924)
- [HIVE-15483](#) - Database and table name is case sensitive when used in show grant
- [HIVE-15761](#) - ObjectStore.getNextNotification could return an empty NotificationEventResponse causing TProtocolException
- [HIVE-15879](#) - Fix HiveMetaStoreChecker.checkPartitionDirs method
- [HIVE-15882](#) - HS2 generating high memory pressure with many partitions and concurrent queries.
- [HIVE-16004](#) - OutOfMemory in SparkReduceRecordHandler with vectorization mode
- [HIVE-16057](#) - SchemaTool ignores --passWord argument if hadoop.security.credential.provider.path is configured
- [HIVE-16060](#) - GenericUDTFJSONTuple's json cache could overgrow beyond its limit
- [HIVE-16079](#) - HS2: high memory pressure due to duplicate Properties objects
- [HIVE-16119](#) - HiveMetaStoreChecker: remove singleThread logic duplication

- [HIVE-16164](#) - Provide mechanism for passing HMS notification ID between transactional and non-transactional listeners.
- [HIVE-16166](#) - HS2 may still waste up to 15% of memory on duplicate strings.
- [HIVE-16196](#) - UDFJson having thread-safety issues
- [HIVE-16286](#) - Log canceled query id
- [HIVE-16308](#) - PreExecutePrinter and PostExecutePrinter should log to INFO level instead of ERROR
- [HIVE-16315](#) - Describe table doesn't show num of partitions
- [HIVE-16334](#) - Query lock contains the query string, which can cause OOM on ZooKeeper
- [HIVE-16363](#) - QueryLifeTimeHooks should catch parse exceptions
- [HIVE-16389](#) - Allow HookContext to access SQLOperationDisplay
- [HIVE-16413](#) - Create table as select does not check ownership of the location
- [HIVE-16422](#) - Should kill running Spark Jobs when a query is cancelled
- [HIVE-16426](#) - Query cancel: improve the way to handle files.
- [HIVE-16430](#) - Add log to show the cancelled query id when cancelOperation is called
- [HIVE-16449](#) - BeeLineDriver should handle query result sorting
- [HIVE-16451](#) - Race condition between HiveStatement.getQueryLog and HiveStatement.runAsyncOnServer
- [HIVE-16487](#) - Serious Zookeeper exception is logged when a race condition happens
- [HIVE-16555](#) - Add a new thrift API call for get_metastore_uuid
- [HIVE-16556](#) - Hive Metastore schema changes to add METASTORE_DB_PROPERTIES table
- [HIVE-16584](#) - Warning messages should use LogHelper.println instead of printing to the infoStream directly
- [HIVE-16593](#) - SparkClientFactory.stop may prevent JVM from exiting
- [HIVE-16647](#) - Improve the validation output to make the output to stderr and stdout more consistent
- [HIVE-16660](#) - Not able to add partition for views in hive when sentry is enabled
- [HIVE-16665](#) - Race condition in Utilities.GetInputPathsCallable --> createDummyFileForEmptyPartition
- [HIVE-16693](#) - beeline "source" command freezes if you have a comment in it?
- [HIVE-16697](#) - Schema table validator should return a sorted list of missing tables
- [HIVE-16700](#) - Log ZK discovery info (hostname & port) for HTTP mode when connection is established
- [HIVE-16711](#) - Remove property_id column from metastore_db_properties table
- [HIVE-16723](#) - Enable configurable MetaStoreSchemaInfo
- [HIVE-16745](#) - Syntax error in 041-HIVE-16556.mysql.sql script
- [HIVE-16771](#) - Schematool should use MetastoreSchemaInfo to get the metastore schema version from database
- [HUE-7](#) - [metastore] Redesign create table wizard
- [HUE-1176](#) - [jib] Revamp v2
- [HUE-3264](#) - [assist] Add search to the HDFS panel
- [HUE-3273](#) - [assist] Add solr inner panel
- [HUE-3283](#) - Revamped UI Layout and interface styling
- [HUE-3294](#) - [core] New design for Hue 4
- [HUE-3295](#) - HUE-3294 [core] Create initial design alternative
- [HUE-3698](#) - [editor] Unify app permissions with the available interpreters
- [HUE-3947](#) - [editor] Offer to manually run queries in editor when sample data timeouts
- [HUE-3995](#) - [editor] Add links to menu of configured snippets
- [HUE-4847](#) - [metastore] Convert all the metastore operations to the new editor
- [HUE-4918](#) - HUE-7 [metastore] Refresh table in Impala after creating a table if possible
- [HUE-4963](#) - [metastore] Convert partition browsing to the new editor
- [HUE-5261](#) - [editor] Move rowcount polling fetching to frontend
- [HUE-5419](#) - [oozie] Custom workflows graph are not displayed properly in Oozie Dashboard
- [HUE-5608](#) - [editor] Display table column list if user does not have access to table but has some column level access
- [HUE-5830](#) - [editor] Notion of currently selected statement within editor with multi query
- [HUE-5848](#) - [core] Auto logout window can trigger an infinite loop on /notebook/api/notebook/close

- [HUE-5860](#) - [assist] Left and right panels UX logic
- [HUE-5912](#) - [editor] huePubSub .subscribe should be active only for the current editor
- [HUE-5926](#) - [sentry] hiveChooser seems flaky and does not autocomplete columns
- [HUE-5956](#) - [metastore] Create DB only remembers the name from the Hive chooser
- [HUE-5957](#) - [metadata] Clicking on table name should add DB prefix too in autocomplete
- [HUE-6021](#) - [metastore] Allow nicer formatting of metadata description
- [HUE-6097](#) - [responsive] Inject schedule into context panel
- [HUE-6119](#) - [core] Configure default snippets and apps and validate blacklisting
- [HUE-6218](#) - [editor] Decrease GPU usage
- [HUE-6270](#) - [editor] Integrated scheduler sometimes gets a coordinator UUID that does not exist
- [HUE-6279](#) - [assist] Right click not staying open in assist in FF when clicking on DB the first time
- [HUE-6294](#) - [editor] Opening saved query removes the pagination buttons
- [HUE-6472](#) - [assist] Functions dropdown is hidden under below flex panel
- [HUE-6564](#) - [indexer] New indexes not showing up and remove cache from assist
- [IMPALA-278](#) - Mention DIV arithmetic operator.
- [IMPALA-1427](#) - Improvements to "Unknown disk-ID" warning
- [IMPALA-1670](#), [IMPALA-4141](#) - Support multiple partitions in ALTER TABLE ADD PARTITION
- [IMPALA-1861](#) - Simplify conditionals with constant conditions
- [IMPALA-1972](#), [IMPALA-3882](#) - Fix client_request_state_map_lock_contention
- [IMPALA-2020](#) - Inline big number strings
- [IMPALA-2079](#) - Don't fail when write to scratch dir results in error.
- [IMPALA-2328](#) - Read support for min/max Parquet statistics
- [IMPALA-2518](#) - DROP DATABASE CASCADE removes cache directives of tables
- [IMPALA-2522](#) - Improve the reliability and effectiveness of ETL.
- [IMPALA-2550](#) - Switch to per-query exec rpc
- [IMPALA-2615](#) - warn if Status is ignored
- [IMPALA-3079](#) - Fix sequence file writer
- [IMPALA-3200](#) - Implement suballocator for splitting buffers
- [IMPALA-3202](#) - implement spill-to-disk in new buffer pool
- [IMPALA-3203](#) - Implement scalable buffer recycling in buffer pool.
- [IMPALA-3524](#) - Don't process spilled partitions with 0 probe rows
- [IMPALA-3586](#) - Implement union passthrough
- [IMPALA-3654](#) - Parquet stats filtering for IN predicate
- [IMPALA-3671](#) - Add SCRATCH_LIMIT query option.
- [IMPALA-3742](#) - Partitions and sort INSERTs for Kudu tables
- [IMPALA-3748](#) - minimum buffer requirements in planner
- [IMPALA-3748](#) - add query-wide resource acquisition step
- [IMPALA-3785](#) - Record query handle for invalid handle
- [IMPALA-3794](#) - Workaround for Breakpad ID conflicts
- [IMPALA-3905](#) - Implements HdfsScanner::GetNext() for text scans.
- [IMPALA-3909](#) - Populate min/max statistics in Parquet writer
- [IMPALA-3989](#) - Display skew warning for poorly formatted Parquet files
- [IMPALA-4008](#) - Don't bake fields into generated IR functions of OldHashTable
- [IMPALA-4014](#) - Introduce query-wide execution state.
- [IMPALA-4029](#) - Reduce memory requirements for storing file metadata
- [IMPALA-4036](#) - invalid SQL generated for partitioned table with comment
- [IMPALA-4041](#) - Limit catalog and admission control updates to coordinators
- [IMPALA-4163](#) - Add sortby() query hint
- [IMPALA-4166](#) - Add SORT BY sql clause
- [IMPALA-4341](#) - Add metadata load to planner timeline

- [IMPALA-4351](#), [IMPALA-4353](#) - query generator random profile options for INSERT.
- [IMPALA-4355](#) - random query generator: modify statement execution flow to support DML
- [IMPALA-4359](#) - qgen: add UPSERT support
- [IMPALA-4390](#) - Separate ADD and DROP PARTITION syntax
- [IMPALA-4431](#) - Add audit event log control mechanism to prevent disk overflow
- [IMPALA-4477](#) - Update Kudu to latest commit on master
- [IMPALA-4499](#) - Table name missing from exec summary
- [IMPALA-4546](#) - Fix Moscow timezone conversion after 2014
- [IMPALA-4548](#) - BlockingJoinNode should wait for async build thread
- [IMPALA-4549](#) - consistently treat 9999 as upper bound for timestamp year
- [IMPALA-4588](#) - Enable starting the minicluster when offline
- [IMPALA-4611](#) - Checking perms on S3 files is a very expensive no-op
- [IMPALA-4616](#) - Add missing Kudu column options
- [IMPALA-4617](#) - remove IsConstant() analysis from backend
- [IMPALA-4624](#) - Implement Parquet dictionary filtering
- [IMPALA-4631](#) - avoid DCHECK in PlanFragementExecutor::Close().
- [IMPALA-4640](#) - Fix number of rows displayed by parquet-reader tool
- [IMPALA-4647](#) - fix full data load with ninja
- [IMPALA-4648](#) - remove build_thirdparty.sh
- [IMPALA-4649](#) - add a mechanism to pass flags into make
- [IMPALA-4650](#) - Add Protobuf 2.6.1 to toolchain and as a build dependency
- [IMPALA-4651](#) - Add LibEv to build
- [IMPALA-4652](#) - Add crcutil to build
- [IMPALA-4653](#) - fix sticky config variable problem
- [IMPALA-4662](#) - Fix NULL literal handling in Kudu IN list predicates
- [IMPALA-4673](#) - Use --local_library_dir for tzdb startup scratch space
- [IMPALA-4674](#) - port BufferedTupleStream to BufferPool
- [IMPALA-4676](#) - remove vestigial references to getBlockStorageLocations() API
- [IMPALA-4678](#) - move query MemTracker into QueryState
- [IMPALA-4684](#) - check-hbase-nodes.py: Build failing on RHEL7 when trying to start HBase.
- [IMPALA-4686](#) - Fix schema output for INT96 columns in parquet-reader tool
- [IMPALA-4689](#) - Fix computation of last active time
- [IMPALA-4701](#) - make distcc work reliably with clang
- [IMPALA-4707](#) - fix use-after-free in QueryExecMgr
- [IMPALA-4710](#) - There is an error in control audit log file size number
- [IMPALA-4711](#) - clarify is_null semantics in udf.h
- [IMPALA-4716](#) - Expr rewrite causes IllegalStateException
- [IMPALA-4731](#), [IMPALA-397](#), [IMPALA-4728](#) - Sorter crash Impalad instance.
- [IMPALA-4734](#) - Set parquet::RowGroup::sorting_columns
- [IMPALA-4738](#) - STDDEV_SAMP should return NULL for single record input
- [IMPALA-4740](#) - Add option to use hdfsPread() for HDFS hedged reads
- [IMPALA-4747](#) - macros should only evaluate their arguments once
- [IMPALA-4748](#) - crash in TmpFileMgr when hitting process mem limit
- [IMPALA-4752](#) - make ObjectPool more efficient
- [IMPALA-4757](#) - addendum: avoid double underscore in name
- [IMPALA-4758](#) - Upgrade gutil to recent Kudu version.
- [IMPALA-4762](#) - RECOVER PARTITIONS should batch partition updates
- [IMPALA-4764](#) - Add Hedged read metrics
- [IMPALA-4788](#) - Use HashSet in RECOVER PARTITIONS duplicate checks
- [IMPALA-4789](#) - Fix slow metadata loading due to inconsistent paths.

- [IMPALA-4792](#) - Fix number of distinct values for a CASE with constant outputs
- [IMPALA-4801](#) - fix heap use after free for MemTracker
- [IMPALA-4815](#) - , [IMPALA-4817](#), [IMPALA-4819](#): Write and Read Parquet Statistics for remaining types
- [IMPALA-4820](#) - avoid writing unencrypted data during write cancellation
- [IMPALA-4828](#) - Add "Known Issues" item for
- [IMPALA-4831](#) - enforce BufferPool reservation invariants
- [IMPALA-4839](#) - Remove implicit 'localhost' for KUDU_MASTER_HOSTS
- [IMPALA-4840](#) - Fix REFRESH performance regression.
- [IMPALA-4846](#) - Upgrade Snappy to 1.1.4
- [IMPALA-4849](#) - IllegalStateException from rewritten CASE expr
- [IMPALA-4858](#) - add more info to MemLimitExceeded errors
- [IMPALA-4859](#) - Push down IS NULL / IS NOT NULL to Kudu
- [IMPALA-4877](#) - fix precedence of unary -/+
- [IMPALA-4880](#) - Clarify synchronization policy for 'done_' in KuduScanNode
- [IMPALA-4883](#) - Union Codegen
- [IMPALA-4884](#) - Add JVM heap and non-heap usage in metrics and UI
- [IMPALA-4885](#) - Expose Jvm thread info in web UI
- [IMPALA-4890](#) - /5143: Coordinator race involving TearDown()
- [IMPALA-4892](#) - Session ID included in error message
- [IMPALA-4893](#) - Efficiently update the rows read counter for sequence file
- [IMPALA-4897](#) - AnalysisException: specified cache pool does not exist
- [IMPALA-4923](#) - reduce memory transfer for selective scans
- [IMPALA-4926](#) - Upgrade LZ4 to 1.7.5
- [IMPALA-4937](#) - Remove unused kudu scanner keep alive variable
- [IMPALA-4943](#) - Speed up block md loading for add/recover partition calls.
- [IMPALA-4946](#) - fix hang in BufferPool
- [IMPALA-4955](#) - Fix integer overflow in hdfs table size accounting
- [IMPALA-4959](#) - Avoid picking up the system's boost cmake module
- [IMPALA-4962](#) - Fix SHOW COLUMN STATS for HS2
- [IMPALA-4965](#) - Authorize access to runtime profile and exec summary
- [IMPALA-4966](#) - Add flatbuffers to build
- [IMPALA-4983](#) - Set toolchain version to include LZ4 build flags
- [IMPALA-4988](#) - Add query option read_parquet_statistics
- [IMPALA-4996](#) - Single-threaded KuduScanNode
- [IMPALA-5003](#) - Constant propagation in scan conjuncts
- [IMPALA-5008](#) - Fix reading stats for TINYINT and SMALLINT
- [IMPALA-5021](#) - Fix count(*) remaining rows overflow in Parquet.
- [IMPALA-5027](#) - addendum - remove unneeded #define
- [IMPALA-5030](#) - Adds support for NVL2() function
- [IMPALA-5034](#) - Update Breakpad to newer version
- [IMPALA-5038](#) - Fix file size regex to include bytes
- [IMPALA-5041](#) - Allow AuthManager::Init() to be called more than once
- [IMPALA-5042](#) - Use a HashSet instead of ArrayList for O(1) look ups
- [IMPALA-5056](#) - Ensure analysis uses 'fresh' catalog after metadata loading
- [IMPALA-5057](#) - Upgrade glog (0.3.4-p2) and gflags
- [IMPALA-5073](#) - Part 1: add option to use mmap() for buffer pool
- [IMPALA-5077](#) - add NUMA and current cpu to CpuInfo
- [IMPALA-5080](#) - OutOfMemory PermGen space
- [IMPALA-5110](#) - Add deb support to dump_breakpad_symbols.py
- [IMPALA-5111](#) - Fix check when creating NOT NULL PK col in Kudu

- [IMPALA-5113](#) - fix dirty unpinned invariant
- [IMPALA-5120](#) - Default to partitioned join when stats are missing
- [IMPALA-5123](#) - Fix ASAN use after free in timezone_db
- [IMPALA-5125](#) - SimplifyConditionalsRule incorrectly handles aggregates
- [IMPALA-5127](#) - Add history_max option
- [IMPALA-5130](#) - fix race in MemTracker::EnableReservationReporting()
- [IMPALA-5137](#) - Support Kudu UNIXTIME_MICROS as Impala TIMESTAMP.
- [IMPALA-5140](#) - improve docs building guidelines
- [IMPALA-5144](#) - Remove sortby() hint
- [IMPALA-5147](#) - Add the ability to exclude hosts from query execution
- [IMPALA-5154](#) - Handle 'unpartitioned' Kudu tables
- [IMPALA-5158](#) - Part 1: include untracked memory in MemTracker dumps
- [IMPALA-5166](#) - clean up BufferPool counters
- [IMPALA-5167](#) - Reduce the number of Kudu clients created
- [IMPALA-5169](#) - Add support for async pins in buffer pool
- [IMPALA-5171](#) - update RAT excluded files list
- [IMPALA-5172](#) - crash in tcmmalloc::CentralFreeList::FetchFromOneSpans
- [IMPALA-5173](#) - crash with hash join feeding directly into nlj
- [IMPALA-5174](#) - Bump gflags to 2.2.0-p1
- [IMPALA-5180](#) - Don't use non-deterministic exprs in partition pruning
- [IMPALA-5181](#) - Extract PYPI metadata from a webpage
- [IMPALA-5182](#) - Explicitly close connection to impalad on error from shell
- [IMPALA-5184](#) - build fe against both Hive 1 & 2 APIs
- [IMPALA-5187](#) - , IMPALA-5208: Bump Breakpad Version, undo IMPALA-3794
- [IMPALA-5188](#) - Add slot sorting in TupleDescriptor::LayoutEquals()
- [IMPALA-5189](#) - Pin version of setuptools-scm
- [IMPALA-5192](#) - Don't bake MemPool* into IR
- [IMPALA-5197](#) - Erroneous corrupted Parquet file message
- [IMPALA-5198](#) - Error messages are sometimes dropped before reaching client
- [IMPALA-5207](#) - ,IMPALA-5214: distcc fixes
- [IMPALA-5217](#) - KuduTableSink checks null constraints incorrectly
- [IMPALA-5220](#) - memory maintenance cleanup
- [IMPALA-5221](#) - Avoid re-use of stale SASL contexts.
- [IMPALA-5222](#) - don't call Bits::Log2*() functions
- [IMPALA-5223](#) - Add waiting for HBase Zookeeper nodes to retry loop
- [IMPALA-5224](#) - remove defunct codehaus repository
- [IMPALA-5229](#) - huge page-backed buffers with TCMalloc
- [IMPALA-5230](#) - fix non-functional impalad under ASAN
- [IMPALA-5232](#) - Parquet reader error message prints memory address instead of value
- [IMPALA-5235](#) - Initialize resourceProfile_ with a dummy value
- [IMPALA-5238](#) - transfer reservations between trackers
- [IMPALA-5258](#) - Pass CMAKE_BUILD_TYPE to Impala-Izo
- [IMPALA-5259](#) - Add REFRESH FUNCTIONS <db> statement
- [IMPALA-5261](#) - Heap use-after-free in HdfsSequenceTableWriter
- [IMPALA-5273](#) - Replace StringCompare with glibc memcmp
- [IMPALA-5282](#) - Handle overflows in computeResourceProfile().
- [IMPALA-5294](#) - Kudu INSERT partitioning fails with constants
- [IMPALA-5301](#) - Set Kudu minicluster memory limit
- [IMPALA-5304](#) - reduce transfer of Parquet decompression buffers
- [IMPALA-5309](#) - Adds TABLESAMPLE clause for HDFS table refs.

- [IMPALA-5318](#) - Generate access events with fully qualified table names
- [IMPALA-5324](#) - Fix version check in EvalDictionaryFilters
- [IMPALA-5325](#) - Do not update totalHdfsBytes_/numHdfsFiles_ on Catalogd
- [IMPALA-5331](#) - Use new libHDFS API to address "Unknown Error 255"
- [IMPALA-5333](#) - Add support for Impala to work with ADLS
- [IMPALA-5338](#) - Fix Kudu timestamp column default values
- [IMPALA-5339](#) - Fix analysis with sort.columns and expr rewrites
- [IMPALA-5340](#) - Query profile displays stale query state
- [IMPALA-5342](#) - Add comments of loaded tables in the response of GetTables
- [IMPALA-5347](#) - Parquet scanner microoptimizations
- [IMPALA-5355](#) - Fix the order of Sentry roles and privileges
- [IMPALA-5357](#) - Fix unixtime to UTC TimestampValue perf
- [IMPALA-5358](#) - Fix repeatable table sample.
- [IMPALA-5363](#) - Reset probe_batch_ after reaching limit
- [IMPALA-5364](#) - Correct title of query locations table
- [IMPALA-5375](#) - Builds on CentOS 6.4 failing with broken python dependencies
- [IMPALA-5377](#) - Impala may crash if given a fragment instance while restarting
- [IMPALA-5378](#) - Disk IO manager needs to understand ADLS
- [IMPALA-5381](#) - Adds DEFAULT_JOIN_DISTRIBUTION_MODE query option.
- [IMPALA-5383](#) - Fix PARQUET_FILE_SIZE option for ADLS
- [IMPALA-5388](#) - Only retry RPC on lost connection in send call
- [IMPALA-5391](#) - remove C++11 from UDF header
- [IMPALA-5411](#) - Avoid log spew from GetRuntimeProfileStr
- [IMPALA-5419](#) - Check for cancellation when building hash tables
- [IMPALA-5424](#) - Ignore errors when removing minidumps folder
- [IMPALA-5426](#) - Update Hive schema script to 1.1.0
- [IMPALA-5432](#) - Remove invalid DCHECK from SetMemLimitExceeded
- [IMPALA-5438](#) - Always eval union const exprs in subplan.
- [IMPALA-5454](#) - Work around template rendering bug in /memz
- [IMPALA-5469](#) - Fix exception when processing catalog update
- [IMPALA-5487](#) - Fix race in RuntimeProfile::toThrift()
- [IMPALA-5537](#) - Retry RPC on some exceptions with SSL connection
- [IMPALA-5558](#) - /IMPALA-5576: Reopen stale client connection
- [IMPALA-5562](#) - Only recomputeMemLayout() if tuple has a layout.
- [IMPALA-3905](#) - Implements HdfsScanner::GetNext() for LZO text scans.
- [IMPALA-5172](#) - fix incorrect cast in call to LZO decompress
- [OOZIE-2387](#) - Oozie is Unable to handle Spaces in file/archive tag.
- [OOZIE-2844](#) - Increase stability of Oozie actions when log4j.properties is missing or not readable
- [OOZIE-2872](#) - Address backward compatibility issue introduced by OOOZIE-2748
- [OOZIE-2874](#) - Make the Launcher Mapper map-only job's InputFormat class pluggable.
- [OOZIE-2908](#) - Fix typo in oozie.actions.null.args.allowed property in oozie-default.xml
- [OOZIE-2923](#) - Improve Spark options parsing
- [PARQUET-389](#) - Support predicate push down on missing columns.
- [PIG-5110](#) - Removing schema alias and :: coming from parent relation
- [SENTRY-1556](#) - Simplify privilege cleaning.
- [SENTRY-1644](#) - Partition ACLs disappear after renaming Hive table with partitions
- [SENTRY-1646](#) - Unable to truncate table <database>.<tablename>; from "default" databases
- [SENTRY-1683](#) - MetastoreCacheInitializer has a race condition in handling results list
- [SENTRY-1730](#) - Optimized FileInputStream /FileOutputStream
- [SENTRY-1759](#) - UpdatableCache leaks connections

- [SENTRY-1788](#) - Sentry Store may use JDO object after the associated data is removed in DB
- [SOLR-6673](#) - MDC based logging of collection, shard etc.
- [SOLR-8836](#) - Return 400, and a SolrException when an invalid json is provided to the update handler instead of 500.
- [SOLR-9153](#) - Update Apache commons beanutils version to 1.9.2
- [SOLR-9527](#) - Improve distribution of replicas when restoring a collection
- [SOLR-10076](#) - ofHide keystore and truststore passwords from /admin/info/* outputs. - Updated JMXJsonServlet to use new API - Set redaction default true
- [SPARK-16138](#) - Try to cancel executor requests only if we have at least 1
- [SPARK-20922](#) - Unsafe deserialization in Spark LauncherConnection.
- [ZOOKEEPER-2693](#) - for CVE-2017-5637.

Issues Fixed in CDH 5.11.x

The following topics describe issues fixed in CDH 5.11.x, from newest to oldest release. You can also review [What's New in CDH 5.11.x](#) on page 1822 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.11.2

CDH 5.11.2 fixes the following issues:

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.11.2:

- [FLUME-2752](#) - Fix AvroSource startup resource leaks
- [FLUME-2905](#) - Fix NetcatSource file descriptor leak if startup fails
- [HADOOP-12751](#) - While using kerberos Hadoop incorrectly assumes names with '@' to be non-simple .
- [HADOOP-14141](#) - Store KMS SSL keystore password in catalina.properties
- [HADOOP-14242](#) - Make KMS Tomcat SSL property sslEnabledProtocols and clientAuth configurable
- [HADOOP-14511](#) - WritableRpcEngine.Invocation#toString NPE on null parameters
- [HDFS-6757](#) - Simplify lease manager with INodeID
- [HDFS-8856](#) - Make LeaseManager#countPath O(1).
- [HDFS-10220](#) - A large number of expired leases can make namenode unresponsive and cause failover
- [HDFS-10506](#) - OIV's ReverseXML processor cannot reconstruct some snapshot details
- [HDFS-11579](#) - Make HttpFS Tomcat SSL property sslEnabledProtocols and clientAuth configurable
- [HDFS-11708](#) - Positional read will fail if replicas moved to different DNs after stream is opened
- [HDFS-11741](#) - Long running balancer may fail due to expired DataEncryptionKey
- [HDFS-11861](#) - ipc.Client.Connection#sendRpcRequest should log request name
- [HDFS-11881](#) - NameNode consumes a lot of memory for snapshot diff report generation
- [HDFS-11960](#) - Successfully closed files can stay under-replicated
- [HDFS-12042](#) - Lazy initialize AbstractINodeDiffList#diffs for snapshots to reduce memory consumption
- [HDFS-12139](#) - HTTPFS liststatus returns incorrect pathSuffix for path of file
- [YARN-2780](#) - Log aggregated resource allocation in rm-appsummary.log
- [YARN-6368](#) - Decommissioning an NM results in a -1 exit code
- [YARN-6615](#) - AmlpFilter drops query parameters on redirect
- [HBASE-15720](#) - Print row locks at the debug dump page
- [HBASE-15837](#) - Memstore size accounting is wrong if postBatchMutate() throws exception
- [HBASE-16033](#) - Add more details in logging of responseTooSlow/TooLarge
- [HBASE-16630](#) - Fragmentation in long running Bucket Cache
- [HBASE-16739](#) - Timed out exception message should include encoded region name
- [HBASE-16977](#) - VerifyReplication should log a printable representation of the row keys
- [HBASE-17131](#) - Avoid livelock caused by HRegion#processRowsWithLocks
- [HBASE-17501](#) - guard against NPE while reading FileTrailer and HFileBlock

- [HBASE-17587](#) - Do not Rethrow DoNotRetryIOException as UnknownScannerException
- [HBASE-17673](#) - Monitored RPC Handler not shown in the WebUI
- [HBASE-17688](#) - MultiRowRangeFilter not working correctly if given same start and stop RowKey
- [HBASE-17710](#) - HBase in standalone mode creates directories with 777 permission
- [HBASE-17731](#) - Fractional latency reporting in MultiThreadedAction
- [HBASE-17798](#) - RpcServer.Listener.Reader can abort due to CancelledKeyException
- [HBASE-17970](#) - Set yarn.app.mapreduce.am.staging-dir when starting MiniMRCluster
- [HBASE-18096](#) - Limit HFileUtil visibility and add missing annotations
- [HIVE-9567](#) - JSON SerDe not escaping special chars when writing char/varchar data
- [HIVE-10209](#) - FetchTask with VC may fail because ExecMapper.done is true
- [HIVE-11418](#) - Dropping a database in an encryption zone with CASCADE and trash enabled fails
- [HIVE-11592](#) - ORC metadata section can sometimes exceed protobuf message size limit
- [HIVE-11878](#) - ClassNotFoundException can possibly occur if multiple jars are registered one at a time in Hive
- [HIVE-12274](#) - Increase width of columns used for general configuration in the metastore.
- [HIVE-12551](#) - Fix several kryo exceptions in branch-1
- [HIVE-12762](#) - Common join on parquet tables returns incorrect result when hive.optimize.index.filter set to true
- [HIVE-13330](#) - ORC vectorized string dictionary reader does not differentiate null vs empty string dictionary
- [HIVE-13947](#) - HoS prints wrong number for hash table size in map join scenario
- [HIVE-14178](#) - Hive::needsToCopy should reuse FileUtils::equalsFileSystem
- [HIVE-14564](#) - Column Pruning generates out of order columns in SelectOperator which cause ArrayIndexOutOfBoundsException.
- [HIVE-15122](#) - Hive: Upcasting types should not obscure stats (min/max/ndv)
- [HIVE-16004](#) - OutOfMemory in SparkReduceRecordHandler with vectorization mode
- [HIVE-16060](#) - GenericUDTFJSONTuple's json cache could overgrow beyond its limit
- [HIVE-16291](#) - Hive fails when unions a parquet table with itself
- [HIVE-16413](#) - Create table as select does not check ownership of the location
- [HIVE-16559](#) - Parquet schema evolution for partitioned tables may break if table and partition serdes differ
- [HIVE-16593](#) - SparkClientFactory.stop may prevent JVM from exiting
- [HIVE-16647](#) - Improve the validation output to make the output to stderr and stdout more consistent
- [HIVE-16660](#) - Not able to add partition for views in hive when sentry is enabled
- [HIVE-16665](#) - Race condition in Utilities.GetInputPathsCallable --> createDummyFileForEmptyPartition
- [HIVE-16693](#) - beeline "source" command freezes if you have a comment in it.
- [HIVE-16697](#) - Schema table validator should return a sorted list of missing tables
- [HIVE-16869](#) - Hive returns wrong result when predicates on non-existing columns are pushed down to Parquet reader
- [HIVE-16930](#) - HoS should verify the value of Kerberos principal and keytab file before adding them to spark-submit command parameters
- [HIVE-16935](#) - Hive should strip comments from input before choosing which CommandProcessor to run.
- [HIVE-17050](#) - Multiline queries that have comment in middle fail when executed via "beeline -e"
- [HIVE-17052](#) - Remove logging of predicate filters
- [HIVE-17149](#) - Hdfs directory is not cleared if partition creation failed on HMS
- [HUE-5504](#) - [oozie] Only use JDBC URL from hive2 action when hardcoded
- [HUE-6370](#) - [jobsub] Attempting to create and save any JobSub fails
- [HUE-6398](#) - [jobsub] Fix loading edit design page
- [HUE-6407](#) - [pig] Play button doesn't come back after killing the running pig job
- [HUE-6446](#) - [oozie] User cant edit shared coordinator or bundle
- [HUE-6604](#) - [oozie] Fix timestamp conversion to server timezone
- [HUE-6694](#) - [aws] Gracefully handle bucket creation error for non-DNS compliant names
- [HUE-6695](#) - [aws] Provide user-friendly message when accessing forbidden paths
- [HUE-6696](#) - [aws] Display user friendly message when accessing non-DNS compliant bucket

- [HUE-6702](#) - [about] About page is accessible without authentication
- [HUE-6710](#) - [notebook] Application reachable directly by users without granted access
- [HUE-6791](#) - [search] Protect against pivot facets conflicting with nested facets
- [HUE-6813](#) - [metastore] Fix database comment last character truncation
- [HUE-6814](#) - [search] Only return distinct usernames in security impersonate dropdown
- [HUE-6819](#) - [oozie] Set generic widget for generic actions in oozie graph
- [HUE-6856](#) - [search] Protect against reflected XSS in search query parameters
- [HUE-6950](#) - [saml] Create home directory for new user login
- [HUE-6995](#) - [oozie] Set minimum width of a workflow node
- [IMPALA-4276](#) - Profile displays non-default query options set by planner
- [IMPALA-4546](#) - Fix Moscow timezone conversion after 2014
- [IMPALA-4631](#) - don't use floating point operations for time unit conversions
- [IMPALA-4716](#) - Expr rewrite causes IllegalStateException
- [IMPALA-4738](#) - STDDEV_SAMP should return NULL for single record input
- [IMPALA-4962](#) - Fix SHOW COLUMN STATS for HS2
- [IMPALA-5021](#) - Fix count(*) remaining rows overflow in Parquet.
- [IMPALA-5056](#) - Ensure analysis uses 'fresh' catalog after metadata loading
- [IMPALA-5154](#) - Handle 'unpartitioned' Kudu tables
- [IMPALA-5172](#) - Buffer overrun for Snappy decompression
- [IMPALA-5187](#) - Bump breakpad version to include the fix for Breakpad #681, re-enable the strict check that was disabled in IMPALA-3794.
- [IMPALA-5189](#) - Pin version of setuptools-scm
- [IMPALA-5197](#) - Erroneous corrupted Parquet file message
- [IMPALA-5198](#) - Error messages are sometimes dropped before reaching client
- [IMPALA-5217](#) - KuduTableSink checks null constraints incorrectly
- [IMPALA-5223](#) - Add waiting for HBase Zookeeper nodes to retry loop
- [IMPALA-5301](#) - Set Kudu minicluster memory limit
- [IMPALA-5318](#) - Generate access events with fully qualified table names
- [IMPALA-5355](#) - Fix the order of Sentry roles and privileges
- [IMPALA-5363](#) - Reset probe_batch_ after reaching limit
- [IMPALA-5419](#) - Check for cancellation when building hash tables
- [IMPALA-5469](#) - Fix exception when processing catalog update
- [IMPALA-5487](#) - Fix race in RuntimeProfile::toThrift()
- [IMPALA-5524](#) - Fixes NPE during planning with DISABLE_UNSAFE_SPILLS=1
- [IMPALA-5554](#) - sorter DCHECK on null column
- [IMPALA-5580](#) - fix Java UDFs that return NULL strings
- [IMPALA-5615](#) - Fix compute incremental stats for general partition exprs
- [IMPALA-5623](#) - Fix lag() on STRING cols to release UDF mem
- [IMPALA-5638](#) - Fix Kudu table set tblproperties inconsistencies
- [IMPALA-5657](#) - Fix a couple of bugs with FunctionCallExpr and IGNORE NULLS
- [IMPALA-5172](#) - fix incorrect cast in call to LZO decompress
- [KITE-1155](#) - Deleting an already deleted empty path should not fail the job
- [OOZIE-2816](#) - Strip out the first command word from Sqoop action if its "sqoop"
- [OOZIE-2844](#) - Increase stability of Oozie actions when log4j.properties is missing or not readable
- [OOZIE-2872](#) - Address backward compatibility issue introduced by OOOZIE-2748
- [OOZIE-2908](#) - Fix typo in oozie.actions.null.args.allowed property in oozie-default.xml
- [OOZIE-2923](#) - Improve Spark options parsing
- [OOZIE-2984](#) - Parse spark-defaults.conf values with spaces without needing the quotes
- [PARQUET-389](#) - Support predicate push down on missing columns.
- [PIG-3567](#) - LogicalPlanPrinter throws OOM for large scripts

- [PIG-3655](#) - BinStorage and InterStorage approach to record markers is broken
- [SENTRY-1644](#) - Partition ACLs disappear after renaming Hive table with partitions.
- [SENTRY-1646](#) - Unable to truncate table <database>.<tablename>; from "default" databases
- [SENTRY-1759](#) - UpdatableCache leaks connections
- [SENTRY-1811](#) - Optimize data structures used in HDFS sync
- [SENTRY-1827](#) - Minimize TPathsDump thrift message used in HDFS sync
- [SOLR-6673](#) - MDC-based logging of collection, shard, etc.
- [SOLR-8836](#) - /update should return BAD REQUEST when invalid JSON provided.
- [SOLR-9153](#) - Update beanutils version to 1.9.2
- [SOLR-9527](#) - Solr RESTORE API doesn't distribute the replicas uniformly
- [SOLR-10076](#) - Hiding keystore and truststore passwords from /admin/info/* outputs.
- [SOLR-10889](#) - Stale zookeeper information is used during failover check.
- [SPARK-13278](#) - Launcher fails to start with JDK 9 EA
- [SPARK-15067](#) - YARN executors are launched with fixed perm gen size
- [SPARK-16845](#) - org.apache.spark.sql.catalyst.expressions.GeneratedClass\$SpecificOrdering `grows beyond 64 KB
- [SPARK-19019](#) - PySpark does not work with Python 3.6.0
- [SPARK-19688](#) - Spark on YARN Credentials File set to different application directory.
- [SPARK-20393](#) - Strengthen Spark to prevent XSS vulnerabilities
- [SPARK-20904](#) - Task failures during shutdown cause problems with preempted executors.
- [SPARK-20922](#) - Unsafe deserialization in Spark LauncherConnection.
- [ZOOKEEPER-1653](#) - zookeeper fails to start because of inconsistent epoch.
- [ZOOKEEPER-2040](#) - Server to log underlying cause of SASL connection problems.

Issues Fixed in CDH 5.11.1

CDH 5.11.1 fixes the following issues:

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.11.1:

- [FLUME-3080](#) - Close failure in HDFS Sink might cause data loss
- [FLUME-3085](#) - HDFS Sink can skip flushing some BucketWriters, might lead to data loss
- [HADOOP-11400](#) - GraphiteSink does not reconnect to Graphite after 'broken pipe'
- [HADOOP-11599](#) - Client#getTimeout should use IPC_CLIENT_PING_DEFAULT when IPC_CLIENT_PING_KEY is not configured
- [HADOOP-12672](#) - RPC timeout should not override IPC ping interval
- [HADOOP-13503](#) - Improve SaslRpcClient failure logging
- [HADOOP-13926](#) - S3Guard: S3AFileSystem::listLocatedStatus() to employ MetadataStore
- [HADOOP-14019](#) - Fix some typos in the s3a docs
- [HADOOP-14028](#) - S3A BlockOutputStreams doesn't delete temporary files in multipart uploads or handle part upload failures
- [HADOOP-14051](#) - S3Guard: link docs from index, fix typos
- [HADOOP-14059](#) - typo in s3a rename(self, subdir) error message
- [HADOOP-14092](#) - Typo in hadoop-aws index.md
- [HADOOP-14104](#) - Client should always ask namenode for kms provider path
- [HADOOP-14144](#) - s3guard: CLI diff non-empty after import on new table
- [HADOOP-14172](#) - S3Guard: import does not import empty directory
- [HADOOP-14195](#) - CredentialProviderFactory\$getProviders is not thread-safe
- [HADOOP-14204](#) - S3A multipart commit failing, "UnsupportedOperationException at java.util.Collections\$UnmodifiableList.sort".
- [HADOOP-14215](#) - DynamoDB client should waitForActive on existing tables
- [HADOOP-14236](#) - S3Guard: S3AFileSystem::rename() should move non-listed sub-directory entries

- [HADOOP-14255](#) - S3A to delete unnecessary fake directory objects in mkdirs()
- [HADOOP-14256](#) - [S3A DOC] Correct the format for "Seoul" example
- [HADOOP-14268](#) - Fix markdown itemization in hadoop-aws documents
- [HADOOP-14282](#) - S3Guard: DynamoDBMetadata::prune() should self interrupt correctly
- [HADOOP-14417](#) - Update default SSL cipher list for KMS
- [HDFS-10715](#) - NPE when applying AvailableSpaceBlockPlacementPolicy
- [HDFS-10797](#) - Revert "Disk usage summary of snapshots causes renamed blocks to get counted twice"
- [HDFS-11499](#) - Decommissioning stuck because of failing recovery
- [HDFS-11515](#) - -du throws ConcurrentModificationException
- [HDFS-11689](#) - New exception thrown by DFSClient%isHDFSEncryptionEnabled broke hacky hive code
- [HDFS-11816](#) - Update default SSL cipher list for HttpFS
- [YARN-3251](#) - Fixed a deadlock in CapacityScheduler when computing absoluteMaxAvailableCapacity in LeafQueue
- [YARN-6042](#) - Dump scheduler and queue state information into FairScheduler DEBUG log.
- [YARN-6360](#) - Prevent FS state dump logger from cramming other log files
- [YARN-6432](#) - FairScheduler: Reserve preempted resources for corresponding applications.
- [YARN-6433](#) - Only accessible cgroup mount directories should be selected for a controller.
- [YARN-6453](#) - fairscheduler-statedump.log gets generated regardless of service
- [YARN-6500](#) - Do not mount inaccessible cgroups directories in CgroupsLCEResourcesHandler.
- [HBASE-15125](#) - BackportHBaseFsck's adoptHdfsOrphan function creates region with wrong end key boundary
- [HBASE-15941](#) - HBase repair should not unsplit healthy splitted region
- [HBASE-15955](#) - Disable action in CatalogJanitor#setEnabled should wait for active cleanup scan to finish
- [HBASE-16032](#) - Possible memory leak in StoreScanner
- [HBASE-16238](#) - It's useless to catch SESSIONEXPIRED exception and retry in RecoverableZooKeeper
- [HBASE-16350](#) - Undo server abort from HBASE-14968
- [HBASE-16429](#) - FSHLog: deadlock if rollWriter called when ring buffer filled with appends
- [HBASE-16663](#) - JMX ConnectorServer stopped when unauthorized user try to stop HM/RS/cluster
- [HBASE-16721](#) - Concurrency issue in WAL unflushed seqId tracking
- [HBASE-17265](#) - Region left unassigned in master failover when region failed to open
- [HBASE-17328](#) - Properly dispose of looped replication peers
- [HBASE-17460](#) - enable_table_replication can not perform cyclic replication of a table
- [HBASE-17717](#) - Explicitly use "sasl" ACL scheme for hbase superuser
- [HBASE-17779](#) - disable_table_replication returns misleading message and does not turn off replication
- [HBASE-17792](#) - Use a shared thread pool for AtomicityWriter, AtomicGetReader, AtomicScanReader's connections in TestAcidGuarantees
- [HIVE-11141](#) - Improve RuleRegExp when the Expression node stack gets huge
- [HIVE-11428](#) - Performance: Struct IN() clauses are extremely slow
- [HIVE-11671](#) - Optimize RuleRegExp in DPP codepath
- [HIVE-11842](#) - Improve RuleRegExp by caching some internal data structures
- [HIVE-12179](#) - Add option to not add spark-assembly.jar to Hive classpath
- [HIVE-12768](#) - Thread safety: binary sortable serde decimal deserialization
- [HIVE-14210](#) - ExecDriver should call jobclient.close() to trigger cleanup
- [HIVE-14380](#) - Queries on tables with remote HDFS paths fail in "encryption" checks.
- [HIVE-15282](#) - Different modification times are used when an index is built and when its staleness is checked
- [HIVE-15782](#) - query on parquet table returns incorrect result when hive.optimize.index.filter is set to true
- [HIVE-15879](#) - Fix HiveMetaStoreChecker.checkPartitionDirs method
- [HIVE-15997](#) - Resource leaks when query is cancelled
- [HIVE-16024](#) - MCK Repair Requires nonstrict hive.mapred.mode
- [HIVE-16047](#) - Shouldn't try to get KeyProvider unless encryption is enabled
- [HIVE-16156](#) - FileSinkOperator should delete existing output target when renaming
- [HIVE-16175](#) - Possible race condition in InstanceCache

- [HIVE-16205](#) - Improving type safety in Objectstore
- [HIVE-16297](#) - Improving hive logging configuration variables
- [HIVE-16394](#) - HoS does not support queue name change in middle of session
- [HIVE-16459](#) - Forward channelInactive to RpcDispatcher
- [HIVE-16646](#) - Alias in transform ... as clause shouldn't be case sensitive
- [HUE-6109](#) - [core] Remove the restriction on Document2 invalid chars
- [HUE-6144](#) - [oozie] Add generic XSL template to workflow graph parser
- [HUE-6154](#) - [core] Ace paste event shouldn't reset the cursor position
- [HUE-6155](#) - [core] Remove table extender header throttling for Firefox
- [HUE-6158](#) - [autocomplete] The autocompleter eats characters to the right of the cursor on insert
- [HUE-6193](#) - [converter] Retain last_executed time when creating doc2 object
- [HUE-6197](#) - [impala] Fix XSS Vulnerability in the old editors' error messages
- [HUE-6207](#) - [editor] Avoid to always show the horizontal scroll bar when there's no scrolling needed
- [HUE-6208](#) - [core] The scroll left anchor should reset the horizontal scrollbar position
- [HUE-6212](#) - [oozie] Prevent XSS injection in coordinator cron frequency field
- [HUE-6223](#) - [autocomplete] Fix issue where tables from subsequent statements appear in the autocomplete results
- [HUE-6228](#) - [core] Disable touchscreen detection on Nicescroll
- [HUE-6251](#) - [editor] Log warnings and continue on failed bulk delete and copy actions
- [IMPALA-3641](#) - Fix catalogd RPC responses to DROP IF EXISTS.
- [IMPALA-4088](#) - Assign fix values to the minicluster server ports
- [IMPALA-4293](#) - query profile should include error log
- [IMPALA-4544](#) - ASAN should ignore SEGV and leaks
- [IMPALA-4615](#) - Fix create_table.sql command order
- [IMPALA-4733](#) - Change HBase ports to non-ephemeral
- [IMPALA-4787](#) - Optimize APPX_MEDIAN() memory usage
- [IMPALA-4822](#) - Implement dynamic log level changes
- [IMPALA-4899](#) - Fix parquet table writer dictionary leak
- [IMPALA-4902](#) - Copy parameters map in HdfsPartition.toThrift().
- [IMPALA-4998](#) - Fix missing table lock acquisition.
- [IMPALA-5028](#) - Lock table in /catalog_objects endpoint.
- [IMPALA-5055](#) - Fix DCHECK in parquet-column-readers.cc ReadPageHeader()
- [IMPALA-5088](#) - Fix heap buffer overflow
- [IMPALA-5115](#) - Handle status from HdfsTableSink::WriteClusteredRowBatch
- [IMPALA-5145](#) - Do not constant fold null in CastExprs
- [IMPALA-5156](#) - Drop VLOG level passed into Kudu client
- [IMPALA-5186](#) - Handle failed CreateAndOpenScanner() in MT scan.
- [IMPALA-5193](#) - Initialize decompressor before finding first tuple
- [IMPALA-5251](#) - Fix propagation of input exprs' types in 2-phase agg
- [IMPALA-5252](#) - Fix crash in HiveUdfCall::GetStringVal() when mem_limit exceeded
- [IMPALA-5253](#) - Use appropriate transport for StatestoreSubscriber
- [IMPALA-5322](#) - Fix a potential crash in Frontend & Catalog JNI startup
- [OOZIE-2739](#) - Remove property expansion pattern from ShellMain's log4j properties content
- [OOZIE-2818](#) - Can't overwrite oozie.action.max.output.data on a per-workflow basis
- [OOZIE-2819](#) - Make Oozie REST API accept multibyte characters for script Actions
- [SENTRY-1508](#) - Revert ""REVERT: MetastorePlugin.java does not handle properly initialization failure
- [SENTRY-1605](#) - SENTRY-1508 need to be fixed because of Kerberos initialization issue
- [SENTRY-1683](#) - MetastoreCacheInitializer has a race condition in handling results list
- [SENTRY-1714](#) - MetastorePlugin.java should quietly return from renameAuthzObject() when both paths are null
- [SOLR-9836](#) - Add ability to recover from leader when index corruption is detected on SolrCore creation.
- [SOLR-9848](#) - Lower solr.cloud.wait-for-updates-with-stale-state-pause back down from 7 seconds.

- [SOLR-10360](#) - Remove an extra space from Hadoop distcp cmd used by Solr backup/restore
- [SOLR-10430](#) - Add ls command to ZkCLI for listing sub-dirs
- [SPARK-14930](#) - [SPARK-13693] Fix race condition in CheckpointWriter.stop()
- [SPARK-19178](#) - [SQL][Backport-to-1.6] convert string of large numbers to int should return null
- [SPARK-19263](#) - DAGScheduler should avoid sending conflicting task set.
- [SPARK-19537](#) - Move pendingPartitions to ShuffleMapStage.
- [SPARK-20435](#) - [CORE] More thorough redaction of sensitive information
- [SQOOP-3123](#) - Introduce escaping logic for column mapping parameters (same what Sqoop already uses for the DB column names), thus special column names (e.g. containing '#' character) and mappings related to those columns can be in the same format (thus not confusing the end users), and also eliminates the related AVRO format clashing issues.
- [SQOOP-3140](#) - Removing deprecated mapred.map.max.attempts, mapred.reduce.max.attempts entries and using the new constants directly from Hadoop instead
- [SQOOP-3159](#) - Sqoop (export + --table) with Oracle table_name having '\$' fails with error

Issues Fixed in CDH 5.11.0

CDH 5.11.0 fixes the following issues.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.11.0:

- [AVRO-1973](#) - Upgrade netty to the latest 3.x release
- [FLUME-2318](#) - Make SpoolingDirectorySource able to handle empty files
- [FLUME-2514](#) - Speed up TestFileChannelRestart tests
- [FLUME-2567](#) - Remove external Maven repositories
- [FLUME-2593](#) - ResettableFileInputStream returns negate values from read() method
- [FLUME-2619](#) - SpoolDir source should log channel exceptions
- [FLUME-2704](#) - Configurable poll delay for spooling directory source
- [FLUME-2718](#) - HTTP Source to support generic Stream Handler
- [FLUME-2725](#) - HDFS Sink does not use configured timezone for rounding
- [FLUME-2761](#) - Move Hive sink out of preview mode
- [FLUME-2798](#) - Malformed Syslog messages can lead to OutOfMemoryException
- [FLUME-2801](#) - Performance improvement on TailDir source
- [FLUME-2806](#) - flume-ng.ps1 Error running script to start an agent on Windows
- [FLUME-2810](#) - Add static Schema URL to AvroEventSerializer configuration
- [FLUME-2812](#) - Fix semaphore leak causing java.lang.Error: Maximum permit count exceeded in MemoryChannel
- [FLUME-2835](#) - Hive Sink tests need to create table with transactional property set
- [FLUME-2857](#) - Make Kafka Source/Channel/Sink restore default values when live updating config
- [FLUME-2875](#) - Allow RollingFileSink to specify a file prefix and a file extension.
- [FLUME-2881](#) - Windows Launch Script fails in plugins dir code
- [FLUME-2889](#) - Fixes to DateTime computations
- [FLUME-2890](#) - Typo in Twitter source warning
- [FLUME-2911](#) - Add include pattern option in SpoolDir source
- [FLUME-2936](#) - KafkaSource tests: Use loopback instead of host address
- [FLUME-2937](#) - Integrate checkstyle for non-test classes
- [FLUME-2939](#) - Update recursive SpoolDir source to use Java 7 APIs
- [FLUME-2941](#) - Integrate checkstyle for test classes
- [FLUME-2949](#) - checkstyle: Make suppressions compatible with Windows
- [FLUME-2950](#) - ReliableSpoolingFileEventReader.rollCurrentFile is broken
- [FLUME-2952](#) - SyslogAgent: Avoid NPE on stop()
- [FLUME-2955](#) - Add file path to the header in TaildirSource

- [FLUME-2959](#) - Fix issues with flume-checkstyle module
- [FLUME-2966](#) - Fix NPE in JMS Source
- [FLUME-2971](#) - Add secure Kafka Sink/Source/Channel setup to the User Guide
- [FLUME-2974](#) - Fix some broken SpoolDirSource unit tests
- [FLUME-2998](#) - Add a missing configuration parameter to SequenceSource docs
- [FLUME-2999](#) - Kafka channel and sink should enable statically assigned partition per event via header
- [FLUME-3027](#) - Change Kafka Channel to clear offsets map after commit
- [FLUME-3049](#) - Make HDFS sink rotate more reliably in secure mode
- [HADOOP-10823](#) - TestReloadingX509TrustManager is flaky
- [HADOOP-11447](#) - Add a more meaningful toString method to SampleStat and MutableStat.
- [HADOOP-11619](#) - FTPFileSystem should override getDefaultPort.
- [HADOOP-11859](#) - PseudoAuthenticationHandler fails with httpcomponents v4.4
- [HADOOP-11878](#) - FileContext#fixRelativePart should check for not null for a more informative exception.
- [HADOOP-12009](#) - Clarify FileSystem.listStatus() sorting order & fix FileSystemContractBaseTest:testListStatus.
- [HADOOP-12655](#) - TestHttpServer.testBindAddress bind port range is wider than expected.
- [HADOOP-12667](#) - s3a to support createNonRecursive API
- [HADOOP-12774](#) - s3a should use UGI.getCurrentUser.getShortname() for username
- [HADOOP-12977](#) - s3a to handle delete("/", true) robustly
- [HADOOP-13037](#) - Refactor Azure Data Lake Store as an independent FileSystem
- [HADOOP-13119](#) - Add ability to secure log servlet using proxy users. Contribute by Yuanbo Liu.
- [HADOOP-13164](#) - Optimize S3AFileSystem::deleteUnnecessaryFakeDirectories
- [HADOOP-13204](#) - Add support for SSE-KMS and SSE-C in s3a filesystem.
- [HADOOP-13207](#) - Specify FileSystem listStatus, listFiles and Remoteliterator
- [HADOOP-13257](#) - Improve Azure Data Lake contract tests
- [HADOOP-13309](#) - Document S3A known limitations in file ownership and permission model
- [HADOOP-13336](#) - S3A to support per-bucket configuration (S3Guard specific changes)
- [HADOOP-13345](#) - S3Guard: Ensure GenericOptionsParser is used for S3Guard CLI
- [HADOOP-13375](#) - o.a.h.security.TestGroupsCaching.testBackgroundRefreshCounters seems flaky.
- [HADOOP-13433](#) - Race in UGI.reloginFromKeytab
- [HADOOP-13446](#) - Support running isolated unit tests separate from AWS integration tests
- [HADOOP-13448](#) - S3Guard: Define MetadataStore interface.
- [HADOOP-13449](#) - S3Guard: Implement DynamoDBMetadataStore
- [HADOOP-13452](#) - S3Guard: Implement access policy for intra-client consistency with in-memory metadata store
- [HADOOP-13455](#) - S3Guard: Write end user docs, change table autocreate default.
- [HADOOP-13496](#) - Include file lengths in Mismatch in length error for distcp
- [HADOOP-13502](#) - Split fs.contract.is-blobstore flag into more descriptive flags for use by contract tests
- [HADOOP-13508](#) - FsPermission string constructor does not recognize sticky bit
- [HADOOP-13540](#) - improve section on troubleshooting s3a auth problems
- [HADOOP-13560](#) - S3ABlockOutputStream to support huge (many GB) file writes
- [HADOOP-13573](#) - S3Guard: create basic contract tests for MetadataStore implementations
- [HADOOP-13589](#) - S3Guard: Allow execution of all S3A integration tests with S3Guard enabled
- [HADOOP-13599](#) - s3a close() to be non-synchronized, so avoid risk of deadlock on shutdown
- [HADOOP-13614](#) - Purge some superfluous/obsolete S3 FS tests that are slowing test runs down
- [HADOOP-13627](#) - Have an explicit KerberosAuthException for UGI to throw, text from public constants
- [HADOOP-13631](#) - S3Guard: implement move() for LocalMetadataStore, add unit tests
- [HADOOP-13650](#) - S3Guard: Provide command line tools to manipulate metadata store.
- [HADOOP-13651](#) - S3Guard: S3AFileSystem Integration with MetadataStore
- [HADOOP-13674](#) - S3A can provide a more detailed error message when accessing a bucket through an incorrect S3 endpoint
- [HADOOP-13680](#) - fs.s3a.readahead.range to use getLongBytes

- [HADOOP-13692](#) - hadoop-aws should declare explicit dependency on Jackson 2 jars to prevent classpath conflicts
- [HADOOP-13727](#) - S3A: Reduce high number of connections to EC2 Instance Metadata Service caused by InstanceProfileCredentialsProvider
- [HADOOP-13735](#) - ITestS3AFileContextStatistics.testStatistics() failing
- [HADOOP-13742](#) - Expose NumOpenConnectionsPerUser as a metric.
- [HADOOP-13749](#) - KMSClientProvider combined with KeyProviderCache can result in wrong UGI being used
- [HADOOP-13782](#) - Make MutableRates metrics thread-local write, aggregate-on-read
- [HADOOP-13793](#) - S3guard: add inconsistency injection, integration tests
- [HADOOP-13801](#) - regression: ITestS3AMiniYarnCluster failing
- [HADOOP-13804](#) - MutableStat mean loses accuracy if add(long, long) is used
- [HADOOP-13805](#) - UGI.getCurrentUser() fails if user does not have a keytab associated
- [HADOOP-13823](#) - s3a rename: fail if dest file exists
- [HADOOP-13826](#) - S3A Deadlock in multipart copy due to thread pool limits
- [HADOOP-13830](#) - Intermittent failure of ITestS3NContractRootDir#testRecursiveRootListing: "Can not create a Path from an empty string"
- [HADOOP-13847](#) - KMSWebApp should close KeyProviderCryptoExtension
- [HADOOP-13850](#) - s3guard to log choice of metadata store at debug
- [HADOOP-13855](#) - Fix a couple of the s3a statistic names to be consistent with the rest
- [HADOOP-13857](#) - S3AUtils.translateException to map (wrapped) InterruptedExceptions to InterruptedIOEs
- [HADOOP-13876](#) - S3Guard: better support for multi-bucket access
- [HADOOP-13877](#) - S3Guard: fix TestDynamoDBMetadataStore when fs.s3a.s3guard.ddb.table is set
- [HADOOP-13886](#) - s3guard: ITestS3AFileOperationCost.testFakeDirectoryDeletion failure
- [HADOOP-13893](#) - dynamodb dependency -> compile
- [HADOOP-13899](#) - S3Guard: tune dynamodb client & tests
- [HADOOP-13900](#) - Remove snapshot version of SDK dependency from Azure Data Lake Store File System
- [HADOOP-13903](#) - Improvements to KMS logging to help debug authorization errors.
- [HADOOP-13904](#) - DynamoDBMetadataStore to handle DDB throttling
- [HADOOP-13908](#) - S3Guard: Existing tables may not be initialized correctly in DynamoDBMetadataStore
- [HADOOP-13911](#) - Remove TRUSTSTORE_PASSWORD related scripts from KMS
- [HADOOP-13914](#) - S3guard: Improve S3AFileStatus#isEmptyDirectory handling
- [HADOOP-13928](#) - TestAdlFileContextMainOperationsLive.testGetFileContext1 runtime error.
- [HADOOP-13929](#) - ADLS connector should not check in contract-test-options.xml.
- [HADOOP-13931](#) - S3AGuard: Use BatchWriteItem in DynamoDBMetadataStore#put()
- [HADOOP-13934](#) - S3Guard: DynamoDBMetadataStore#move() could be throwing exception due to BatchWriteItem limits
- [HADOOP-13937](#) - S3Guard: Mock bucket locations in MockS3ClientFactory
- [HADOOP-13953](#) - Make FTPFileSystem's data connection mode and transfer mode configurable
- [HADOOP-13956](#) - Read ADLS credentials from Credential Provider.
- [HADOOP-13958](#) - Bump up release year to 2017
- [HADOOP-13960](#) - Initialize DynamoDBMetadataStore without associated S3AFileSystem.
- [HADOOP-13962](#) - Update ADLS SDK to 2.1.4.
- [HADOOP-13976](#) - Path globbing does not match newlines
- [HADOOP-13985](#) - s3guard: add a version marker to every table
- [HADOOP-13988](#) - KMSClientProvider does not work with WebHDFS and Apache Knox w/ProxyUser
- [HADOOP-13995](#) - s3guard cli: make tests easier to run and address failure
- [HADOOP-14001](#) - Improve delegation token validity checking.
- [HADOOP-14003](#) - Make additional KMS tomcat settings configurable
- [HADOOP-14013](#) - S3Guard: fix multi-bucket integration tests
- [HADOOP-14017](#) - User friendly name for ADLS user and group
- [HADOOP-14020](#) - Optimize dirListingUnion

- [HADOOP-14027](#) - S3Guard: Implicitly creating DynamoDB table ignores endpoint config
- [HADOOP-14029](#) - Fix KMSClientProvider for non-secure proxyuser use case
- [HADOOP-14040](#) - Use shaded aws-sdk uber-JAR 1.11.86
- [HADOOP-14041](#) - CLI command to prune old metadata
- [HADOOP-14044](#) - Synchronization issue in delegation token cancel functionality
- [HADOOP-14046](#) - Metastore destruction test creates table w/o version
- [HADOOP-14049](#) - Honour AclBit flag associated to file/folder permission for Azure datalake account
- [HADOOP-14050](#) - Add process name to kms process
- [HADOOP-14079](#) - s3guard: fix breaking link in s3guard.md
- [HADOOP-14083](#) - KMS should support old SSL clients.
- [HADOOP-14085](#) - s3guard: drop unnecessary type assertion and cast.
- [HADOOP-14094](#) - Rethink S3GuardTool options
- [HADOOP-14096](#) - s3guard: regression in dirListingUnion
- [HADOOP-14102](#) - Relax error message assertion in S3A test ITestS3AEncryptionSSEC
- [HADOOP-14110](#) - In S3AFileSystem, make getAmazonClient() package private; export getBucketLocation()
- [HADOOP-14113](#) - Review ADL Docs
- [HADOOP-14114](#) - S3A can no longer handle unencoded + in URIs
- [HADOOP-14125](#) - s3guard tool tests aren't isolated; can't run in parallel
- [HADOOP-14129](#) - ITestS3ACredentialsInURL sometimes fails
- [HADOOP-14130](#) - Simplify DynamoDBClientFactory for creating Amazon DynamoDB clients
- [HADOOP-14131](#) - kms.sh creates bogus dir for tomcat logs
- [HADOOP-14168](#) - S3GuardTool tests should not run if S3Guard is not set up
- [HADOOP-14181](#) - S3Guard: Add validation of DynamoDB region
- [HDFS-3716](#) - Purger should remove stale fsimage ckpt files
- [HDFS-4366](#) - Block Replication Policy Implementation May Skip Higher-Priority Blocks for Lower-Priority Blocks.
- [HDFS-8307](#) - Spurious DNS Queries from hdfs shell
- [HDFS-8640](#) - Make reserved RBW space visible through JMX.
- [HDFS-8883](#) - NameNode Metrics : Add FSNameSystem lock Queue Length
- [HDFS-9145](#) - Tracking methods that hold FSNamesytemLock for too long
- [HDFS-9205](#) - Do not schedule corrupt blocks for replication.
- [HDFS-9390](#) - Block management for maintenance states.
- [HDFS-9391](#) - Update webUI/JMX to display maintenance state info.
- [HDFS-9467](#) - Fix data race accessing writeLockHeldTimeStamp in FSNamesystem
- [HDFS-10292](#) - Add block id when client got Unable to close file exception
- [HDFS-10534](#) - NameNode WebUI should display DataNode usage histogram
- [HDFS-10559](#) - DiskBalancer: Use SHA1 for Plan ID
- [HDFS-10713](#) - Throttle FsNameSystem lock warnings
- [HDFS-10733](#) - NameNode terminated after full GC thinking QJM is unresponsive
- [HDFS-10798](#) - Make the threshold of reporting FSNamesystem lock contention configurable
- [HDFS-10808](#) - DiskBalancer does not execute multi-steps plan-redux
- [HDFS-10817](#) - Add Logging for Long-held NN Read Locks
- [HDFS-10896](#) - Move lock logging logic from FSNamesystem into FSNamesystemLock
- [HDFS-10941](#) - Improve BlockManager#processMisReplicatesAsync log
- [HDFS-10966](#) - Enhance Dispatcher logic on deciding when to give up a source DataNode
- [HDFS-11087](#) - NamenodeFsck should check if the output writer is still writable
- [HDFS-11132](#) - Allow AccessControlException in contract tests when getFileStatus on subdirectory of existing files
- [HDFS-11160](#) - VolumeScanner reports write-in-progress replicas as corrupt incorrectly
- [HDFS-11180](#) - Intermittent deadlock in NameNode when failover happens.
- [HDFS-11197](#) - Listing encryption zones fails when deleting a EZ that is on a snapshotted directory
- [HDFS-11259](#) - Update fsck to display maintenance state info.

- [HDFS-11263](#) - ClassCastException when we use Bzipcodec for Fimage compression
- [HDFS-11265](#) - Extend visualization for Maintenance Mode under Datanode tab in the NameNode UI.
- [HDFS-11271](#) - Typo in NameNode UI
- [HDFS-11275](#) - Check groupEntryIndex and throw a helpful exception on failures when removing ACL.
- [HDFS-11280](#) - Allow WebHDFS to reuse HTTP connections to NN
- [HDFS-11292](#) - log lastWrittenTxId etc info in logSyncAll
- [HDFS-11296](#) - Maintenance state expiry should be an epoch time and not jvm monotonic.
- [HDFS-11306](#) - Print remaining edit logs from buffer if edit log can't be rolled
- [HDFS-11363](#) - Need more diagnosis info when seeing Slow waitForAckedSeqno.
- [HDFS-11378](#) - Verify multiple DataNodes can be decommissioned/maintenance at the same time.
- [HDFS-11379](#) - DFSInputStream may infinite loop requesting block locations
- [HDFS-11390](#) - Add process name to httpfs process
- [HDFS-11411](#) - Avoid OutOfMemoryError in TestMaintenanceState test runs.
- [HDFS-11412](#) - Maintenance minimum replication config value allowable range should be [0, DefaultReplication].
- [HDFS-11418](#) - HttpFS should support old SSL clients
- [HDFS-11441](#) - Add escaping to error message in KMS web UI
- [HDFS-11466](#) - Change dfs.namenode.write-lock-reporting-threshold-ms default from 1000ms to 5000ms
- [MAPREDUCE-5155](#) - Race condition in test case TestFetchFailure cause it to fail
- [MAPREDUCE-5335](#) - Rename Job Tracker terminology in ShuffleSchedulerImpl
- [MAPREDUCE-5883](#) - "Total megabyte-seconds" in job counters is slightly misleading
- [MAPREDUCE-5981](#) - Log levels of certain MR logs can be changed to DEBUG
- [MAPREDUCE-6172](#) - TestDbClasses timeouts are too aggressive
- [MAPREDUCE-6571](#) - JobEndNotification info logs are missing in AM container syslog
- [MAPREDUCE-6639](#) - Process hangs in LocatedFileStatusFetcher if FileSystem.get throws.
- [MAPREDUCE-6715](#) - Fix Several Unsafe Practices
- [MAPREDUCE-6817](#) - The format of job start time in JHS is different from those of submit and finish time.
- [MAPREDUCE-6829](#) - Add peak memory usage counter for each task.
- [YARN-1047](#) - Expose # of pre-emptions as a queue counter
- [YARN-2301](#) - Improved yarn container command
- [YARN-2306](#) - Add test for leakage of reservation metrics in fair scheduler.
- [YARN-2336](#) - Fair scheduler's REST API returns a missing '[' bracket JSON for deep queue tree
- [YARN-2902](#) - Addendum. Killing a container that is localizing can orphan resources in the DOWNLOADING state
- [YARN-2988](#) - Graph#save() may leak file descriptors.
- [YARN-3269](#) - Yarn.nodemanager.remote-app-log-dir could not be configured to fully qualified path
- [YARN-3396](#) - Handle URISyntaxException in ResourceLocalizationService.
- [YARN-3933](#) - FairScheduler: Multiple calls to completedContainer are not safe.
- [YARN-3957](#) - FairScheduler NPE In FairSchedulerQueueInfo causing scheduler page to return 500.
- [YARN-4232](#) - TopCLI console support for HA mode
- [YARN-4363](#) - In TestFairScheduler, testcase should not create FairScheduler redundantly.
- [YARN-4544](#) - All the log messages about rolling monitoring interval are shown with WARN level
- [YARN-4555](#) - TestDefaultContainerExecutor#testContainerLaunchError fails on non-english locale environment
- [YARN-4719](#) - Add a helper library to maintain node state and allows common queries.
- [YARN-4752](#) - Improved preemption in FairScheduler.
- [YARN-4805](#) - Don't go through all schedulers in ParameterizedTestBase.
- [YARN-4882](#) - Change the log level to DEBUG for recovering completed applications
- [YARN-5136](#) - Error in handling event type APP_ATTEMPT_REMOVED to the scheduler
- [YARN-5181](#) - ClusterNodeTracker: add method to get list of nodes matching a specific resourceName.
- [YARN-5182](#) - MockNodes.newNodes creates one more node per rack than requested.
- [YARN-5257](#) - Fix unreleased resources and null dereferences
- [YARN-5308](#) - FairScheduler: Move continuous scheduling related tests to TestContinuousScheduling

- [YARN-5752](#) - TestLocalResourcesTrackerImpl#testLocalResourceCache times out
- [YARN-5798](#) - Set UncaughtExceptionHandler for all FairScheduler threads.
- [YARN-5830](#) - FairScheduler: Avoid preempting AM containers.
- [YARN-5859](#) - TestResourceLocalizationService#testParallelDownloadAttemptsForPublicResource sometimes fails
- [YARN-5890](#) - FairScheduler should log information about AM-resource-usage and max-AM-share for queues
- [YARN-5920](#) - Fix deadlock in TestRMHA.testTransitionedToStandbyShouldNotHang
- [YARN-6061](#) - Add an UncaughtExceptionHandler for critical threads in RM.
- [YARN-6112](#) - UpdateCallDuration is calculated only when debug logging is enabled.
- [YARN-6125](#) - The application attempt's diagnostic message should have a maximum size
- [YARN-6131](#) - FairScheduler: Lower update interval for faster tests.
- [YARN-6144](#) - FairScheduler: preempted resources can become negative.
- [YARN-6163](#) - FS Preemption is a trickle for severely starved applications.
- [YARN-6171](#) - ConcurrentModificationException on FSAppAttempt.containersToPreempt.
- [YARN-6172](#) - FSLeafQueue demand update needs to be atomic.
- [YARN-6193](#) - FairScheduler might not trigger preemption when using DRF.
- [YARN-6194](#) - Cluster capacity in SchedulingPolicy is updated only on allocation file reload.
- [YARN-6210](#) - FairScheduler: Node reservations can interfere with preemption.
- [YARN-6215](#) - Custom Cloudera version. FairScheduler preemption and update should not run concurrently.
- [YARN-6218](#) - Fix TestAMRMClient when using FairScheduler.
- [YARN-6264](#) - AM not launched when a single vcore is available on the cluster.
- [YARN-6448](#) - Continuous scheduling thread crashes while sorting nodes.
- [HBASE-6721](#) - RegionServer Group based Assignment. Adds new rsgroups feature which allows isol
- [HBASE-12949](#) - Scanner can be stuck in infinite loop if the HFile is corrupted
- [HBASE-15295](#) - MutateTableAccess.multiMutate() does not get high priority causing a deadlock
- [HBASE-15315](#) - Remove always set super user call as high priority
- [HBASE-15322](#) - Operations using Unsafe path broken for platforms not having sun.misc.Unsafe.
- [HBASE-15325](#) - ResultScanner allowing partial result will miss the rest of the row if the region is moved between two rpc requests
- [HBASE-15328](#) - sanity check the redirect used to send master info requests to the embedded regionserver.
- [HBASE-15378](#) - Scanner cannot handle heartbeat message with no results
- [HBASE-15435](#) - Add WAL (in bytes) written metric
- [HBASE-15441](#) - Fix WAL splitting when region has moved multiple times
- [HBASE-15587](#) - FSTableDescriptors.getDescriptor() logs stack trace erroneously
- [HBASE-15645](#) - hbase.rpc.timeout is not used in operations of HTable
- [HBASE-15931](#) - Add log for long-running tasks in AsyncProcess
- [HBASE-16062](#) - Improper error handling in WAL Reader/Writer creation
- [HBASE-16165](#) - Decrease RpcServer.callQueueSize before writeResponse causes OOM
- [HBASE-16237](#) - Blocks for hbase:meta table are not cached in L1 cache
- [HBASE-16266](#) - Do not throw ScannerTimeoutException when catch UnknownScannerException
- [HBASE-16304](#) - HRegion#RegionScannerImpl#handleFileNotFoundException may lead to deadlock when trying to obtain write lock on updatesLock
- [HBASE-16460](#) - Can't rebuild the BucketAllocator's data structures when BucketCache uses FileIOEngine
- [HBASE-16527](#) - IOExceptions from DFS client still can cause CatalogJanitor to delete referenced files
- [HBASE-16604](#) - Scanner retries on IOException can cause the scans to miss data
- [HBASE-16649](#) - Truncate table with splits preserved can cause both data loss and truncated data appeared again
- [HBASE-16662](#) - Fix open POODLE vulnerabilities
- [HBASE-16754](#) - All WALSplitter OutputSinks should process compaction events
- [HBASE-16765](#) - backportNew SteppingRegionSplitPolicy, avoid too aggressive spread of regions for small tables. Includesand the amendmentAmend: Make SteppingSplitPolicy the default.
- [HBASE-16807](#) - , RegionServer will fail to report new active Hmaster until HMaster/RegionServer failover.

- [HBASE-16841](#) - Data loss in MOB files after cloning a snapshot and deleting that snapshot
- [HBASE-16931](#) - Setting cell's seqId to zero in compaction flow might cause RS down.
- [HBASE-16960](#) - RegionServer hang when aborting
- [HBASE-16981](#) - Expand Mob Compaction Partition policy from daily to weekly, monthly
- [HBASE-17020](#) - keylen in midkey() dont computed correctly
- [HBASE-17023](#) - Region left unassigned due to AM and SSH each thinking others would do the assignment work
- [HBASE-17044](#) - Fix merge failed before creating merged region leaves meta inconsistent
- [HBASE-17069](#) - RegionServer writes invalid META entries for split daughters in some circumstances
- [HBASE-17206](#) - FSHLog may roll a new writer successfully with unflushed entries
- [HBASE-17238](#) - Wrong in-memory hbase:meta location causing SSH failure
- [HBASE-17241](#) - Avoid compacting already compacted mob files with _del files
- [HBASE-17243](#) - Reuse CompactionPartitionId and avoid creating MobFileName in PartitionedMobCompactor to avoid unnecessary new objects
- [HBASE-17243](#) - Reuse CompactionPartitionId and avoid creating MobFileName in PartitionedMobCompactor to avoid unnecessary new objects
- [HBASE-17275](#) - Assign timeout may cause region to be unassigned forever
- [HBASE-17309](#) - Fix connection leaks in TestAcidGuarantees
- [HBASE-17381](#) - ReplicationSourceWorkerThread can die due to unhandled exceptions
- [HBASE-17409](#) - Limit jsonp callback name to prevent xss
- [HBASE-17437](#) - Support specifying a WAL directory outside of the root directory
- [HBASE-17452](#) - Failed taking snapshot - region Manifest proto-message too large
- [HBASE-17522](#) - Handle JVM throwing runtime exceptions when we ask for details on heap usage the same as a correctly returned 'undefined'.
- [HBASE-17558](#) - ZK dumping jsp should escape HTML.
- [HBASE-17561](#) - table status page should escape values that may contain arbitrary characters.
- [HBASE-17624](#) - ADDENDUMincluding missing parts of TestWALObserver changes from
- [HBASE-17640](#) - Unittest error in TestMobCompactor with different timezone
- [HBASE-17653](#) - [HBASE-17624](#) rsgroup synchronizations will (distributed) deadlock
- [HBASE-17654](#) - RSGroup refactoring.
- [HBASE-17675](#) - ReplicationEndpoint should choose new sinks if a SaslException occurs
- [HBASE-17677](#) - ServerName parsing from directory name should be more robust to errors from guava's HostAndPort
- [HIVE-7723](#) - Explain plan for complex query with lots of partitions is slow due to in-efficient collection used to find a matching ReadEntity
- [HIVE-9138](#) - Add some explain to PTF operator
- [HIVE-9481](#) - allow column list specification in INSERT statement
- [HIVE-10641](#) - create CRC32 UDF
- [HIVE-11147](#) - MetaTool doesn't update FS root location for partitions with space in name
- [HIVE-11293](#) - HiveConnection.setAutoCommit(true) throws exception
- [HIVE-11376](#) - Removes legacy code wrt skipping files with codec to HiveInputFormat from CombineHiveInputFormat
- [HIVE-11594](#) - Analyze Table for column names with embedded spaces
- [HIVE-11849](#) - NPE in HiveHBaseTableShapshotInputFormat in query with just count(*)
- [HIVE-11882](#) - Fetch optimizer should stop source files traversal once it exceeds the hive.fetch.task.conversion.threshold
- [HIVE-11920](#) - ADD JAR failing with URL schemes other than file/ivy/hdfs
- [HIVE-12222](#) - Define port range in property for RPCServer
- [HIVE-12236](#) - Enable SimpleFetchOptimizer for more query types
- [HIVE-12261](#) - schematool version info exit status should depend on compatibility, not equality
- [HIVE-12349](#) - NPE in ORC SARG for IS NULL queries on Timestamp and Date columns
- [HIVE-12465](#) - Hive might produce wrong results when (outer) joins are merged
- [HIVE-12780](#) - Fix the output of the history command in Beeline

- [HIVE-12789](#) - Fix output twice in the history command of Beeline
- [HIVE-12897](#) - Improve dynamic partition loading
- [HIVE-12976](#) - MetaStoreDirectSql doesn't batch IN lists in all cases
- [HIVE-13149](#) - Remove some unnecessary HMS connections from HS2
- [HIVE-13179](#) - Allow custom HiveConf to be passed to Authentication Providers
- [HIVE-13240](#) - GroupByOperator: Drop the hash aggregates when closing operator
- [HIVE-13463](#) - Fix ImportSemanticAnalyzer to allow for different src/dst filesystems
- [HIVE-13510](#) - Dynamic partitioning doesn't work when remote metastore is used
- [HIVE-13726](#) - Improve dynamic partition loading VI
- [HIVE-13864](#) - Beeline ignores the command that follows a semicolon and comment
- [HIVE-13866](#) - flatten callstack for directSQL errors
- [HIVE-13895](#) - HoS start-up overhead in yarn-client mode
- [HIVE-13901](#) - Hivemetastore add partitions can be slow depending on filesystems
- [HIVE-13902](#) - [Refactor] Minimize metastore jar dependencies on task nodes
- [HIVE-13904](#) - Ignore case when retrieving ColumnInfo from RowResolver
- [HIVE-13966](#) - DbNotificationListener: can loose DDL operation notifications
- [HIVE-14204](#) - Optimize loading dynamic partitions
- [HIVE-14423](#) - S3: Fetching partition sizes from FS can be expensive when stats are not available in metastore
- [HIVE-14511](#) - Improve MSCK for partitioned table to deal with special cases Updated changed q.out
- [HIVE-14693](#) - Some partitions will be left out when partition number is the multiple of the option hive.msck.repair.batch.size
- [HIVE-14713](#) - LDAP Authentication Provider should be covered with unit tests
- [HIVE-14764](#) - Enabling "hive.metastore.metrics.enabled" throws OOM in HiveMetastore
- [HIVE-14803](#) - S3: Stats gathering for insert queries can be expensive for partitioned dataset
- [HIVE-14819](#) - FunctionInfo for permanent functions shows TEMPORARY FunctionType
- [HIVE-14820](#) - RPC server for spark inside HS2 is not getting server address properly
- [HIVE-14873](#) - Add UDF for extraction of 'day of week'
- [HIVE-14920](#) - S3: Optimize SimpleFetchOptimizer::checkThreshold()
- [HIVE-15025](#) - Secure-Socket-Layer (SSL) support for HMS
- [HIVE-15060](#) - Remove the autoCommit warning from beeline
- [HIVE-15065](#) - SimpleFetchOptimizer should decide based on metastore stats when available
- [HIVE-15137](#) - metastore add partitions background thread should use current username
- [HIVE-15207](#) - Implement a capability to detect incorrect sequence numbers
- [HIVE-15072](#) - Schematool should recognize missing tables in metastore
- [HIVE-15073](#) - Schematool should detect malformed URIs
- [HIVE-12261](#) - Schematool version info exit status should depend on compatibility, not equality
- [HIVE-15074](#) - Schematool provides a way to detect invalid entries in VERSION table
- [HIVE-15263](#) - Detect the values for incorrect NULL values
- [HIVE-15391](#) - Location validation for table should ignore the values for view
- [HIVE-15392](#) - Refactoring the validate function of HiveSchemaTool to make the output consistent
- [HIVE-15232](#) - Add notification events for functions and indexes
- [HIVE-15338](#) - Wrong result from non-vectorized DATEDIFF with scalar parameter of type DATE/TIMESTAMP
- [HIVE-15346](#) - "values temp table" should not be an input
- [HIVE-15359](#) - CDH=48042skip.footer.line.count doesnt work properly for certain situations
- [HIVE-15410](#) - WebHCat supports get/set table property with its name containing period and hyphen
- [HIVE-15430](#) - Change SchemaTool table validator to test based on the dbType
- [HIVE-15485](#) - Investigate the DoAs failure in HoS
- [HIVE-15517](#) - NOT (x <=> y) returns NULL if x or y is NULL
- [HIVE-15530](#) - Optimize the column stats update logic in table alteration
- [HIVE-15546](#) - Optimize Utilities.getInputPaths() so each listStatus of a partition is done in parallel

- [HIVE-15550](#) - fix arglist logging in schematool
- [HIVE-15551](#) - memory leak in directsql for mysql+bonecp specific initialization
- [HIVE-15572](#) - Improve the response time for query canceling when it happens during acquiring locks
- [HIVE-15735](#) - In some cases, view objects inside a view do not have parents.
- [HIVE-15754](#) - exchange partition is not generating notifications
- [HIVE-15766](#) - DBNotificationlistener leaks JDOPersistenceManager : Use same PersistenceManager for metadata and notification
- [HIVE-15778](#) - DROP INDEX (non-existent) throws NPE when using DbNotificationListener
- [HIVE-15803](#) - msck can hang when nested partitions are present
- [HIVE-15805](#) - Some minor improvement on the validation tool
- [HIVE-15815](#) - Allow to pass some Oozie properties to Spark in HoS
- [HIVE-15872](#) - The PERCENTILE_APPROX UDAF does not work with empty set
- [HIVE-15879](#) - Fix HiveMetaStoreChecker.checkPartitionDirs method
- [HIVE-15881](#) - Use hive.exec.input.listing.max.threads variable name instead of mapred.dfsclient.parallelism.max
- [HIVE-15895](#) - Use HDFS for stats collection temp dir on blob storage
- [HIVE-16014](#) - HiveMetastoreChecker should use hive.metastore.fshandler.threads instead of hive.mv.files.thread for pool size
- [HIVE-16019](#) - Query fails when group by/order by on same column with uppercase name
- [HIVE-16086](#) - Fix HiveMetaStoreChecker.checkPartitionDirsSingleThreaded method
- [HIVE-16176](#) - Fixing the non-zero exit code issue with HiveSchemaTool
- [HUE-5167](#) - Improve SASL error reporting in HUE
- [HUE-5482](#) - [home] Handle multiple home/trash directories by merging them into one
- [HUE-5594](#) - [editor] The new autocomplete dropdown should include popular group by and order by clauses
- [HUE-5659](#) - [home] Ignore history dependencies when importing document from different cluster
- [HUE-5701](#) - [security] Remove last references to \$.cookie and migrate to \$.totalStorage from FB
- [IMPALA-1430](#) - IMPALA-4878,IMPALA-4879: codegen native UDAs
- [IMPALA-2020](#) - 4915, 4936: Add rounding for decimal casts
- [IMPALA-2020](#) - IMPALA-4809: Codegen support for DECIMAL_V2
- [IMPALA-2605](#) - Omit the sort and mini stress tests
- [IMPALA-4055](#) - Speed up to_date() with custom implementation.
- [IMPALA-4263](#) - Fix wrong omission of agg/analytic hash exchanges.
- [IMPALA-4282](#) - Remove max length check for type strings.
- [IMPALA-4370](#) - Divide and modulo result types for DECIMAL version V2
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog
- [IMPALA-4675](#) - Case-insensitive matching of Parquet fields.
- [IMPALA-4702](#) - Fix command line help for webserver_private_key_file
- [IMPALA-4705](#) - IMPALA-4779, IMPALA-4780: Fix some Expr bugs with codegen
- [IMPALA-4725](#) - Query option to control Parquet array resolution.
- [IMPALA-4729](#) - Implement REPLACE()
- [IMPALA-4742](#) - Change "{}".format() to "{0}".format() for Py 2.6
- [IMPALA-4749](#) - hit DCHECK in sorter with scratch limit
- [IMPALA-4767](#) - Workaround for HIVE-15653 to preserve table stats.
- [IMPALA-4808](#) - old hash join can reference invalid memory
- [IMPALA-4809](#) - Enable support for DECIMAL_V2 in decimal_casting.py
- [IMPALA-4810](#) - fix incorrect expr-test decimal types
- [IMPALA-4810](#) - Make DECIMAL expr-test cases table driven
- [IMPALA-4810](#) - Add DECIMAL_V2 query option
- [IMPALA-4813](#) - Round on divide and multiply
- [IMPALA-4821](#) - Update AVG() for DECIMAL_V2
- [IMPALA-4828](#) - Alter Kudu schema outside Impala may crash on read

- [IMPALA-4854](#) - Fix incremental stats with complex types.
- [IMPALA-4916](#) - Fix maintenance of set of item sets in DisjointSet.
- [IMPALA-4929](#) - Safe concurrent access to IR function call graph
- [IMPALA-4933](#) - IMPALA-4931: Simplify SSL initialization on startup
- [IMPALA-4934](#) - Disable Kudu OpenSSL initialization
- [IMPALA-4981](#) - Re-enable spilling with MT_DOP.
- [IMPALA-4995](#) - Fix integer overflow in TopNNode::PrepareForOutput
- [IMPALA-4997](#) - Fix overflows in Sorter::Tupeliterator
- [IMPALA-5005](#) - Don't allow server to send SASL COMPLETE msg out of order
- [IMPALA-5025](#) - Update binutils to 2.26.1
- [IMPALA-5027](#) - make udf headers buildable externally
- [OOZIE-2243](#) - Kill Command does not kill the child job for java action
- [OOZIE-2519](#) - Oozie HA with SSL info is slightly incorrect
- [OOZIE-2584](#) - Eliminate Thread.sleep() calls in TestMemoryLocks
- [OOZIE-2659](#) - TestPauseTransitService is flaky
- [OOZIE-2720](#) - Test failure - TestCoordMaterializeTriggerService#testMaxMatThrottleNotPicked
- [OOZIE-2742](#) - Unable to kill applications based on tag
- [OOZIE-2748](#) - NPE in LauncherMapper.printArgs()
- [OOZIE-2757](#) - Malformed xml is Spark action doc page
- [OOZIE-2777](#) - Config-default.xml longer than 64k results in java.io.UTFDataFormatException
- [OOZIE-2786](#) - Pass Oozie workflow ID and settings to Spark application configuration
- [OOZIE-2787](#) - Oozie distributes application jar twice making the spark job fail
- [OOZIE-2792](#) - Hive2 action is not parsing Spark application ID from log file properly when Hive is on Spark
- [OOZIE-2802](#) - Spark action failure on Spark 2.1.0 due to duplicate sharelibs
- [OOZIE-2803](#) - Mask passwords when printing out configs/args in MapReduceMain and SparkMain
- [OOZIE-2811](#) - Add support for filtering out properties from SparkConfigurationService
- [SENTRY-1476](#) - SentryStore is subject to JDQL injection. Sentry is subject to JDOQL injection
- [SENTRY-1504](#) - NPE in log4j.properties parsingSentryMain trips over missing log4j property
- [SENTRY-1505](#) - CommitContext isn't used by anything and should be removed
- [SENTRY-1507](#) - Sentry should use Datanucleus version of javax.jdo
- [SENTRY-1508](#) - MetastorePlugin.java does not handle properly initialization failure
- [SENTRY-1515](#) - Cleanup exception handling in SentryStore
- [SENTRY-1517](#) - SentryStore should actually use function getMSentryRole to get roles
- [SENTRY-1518](#) - Add metrics for SentryStore transactions
- [SENTRY-1520](#) - Provide mechanism for triggering HMS full snapshot
- [SENTRY-1533](#) - Sentry console metrics reporting interval should be configurable
- [SENTRY-1534](#) - Oracle supports serializable instead of repeatable-read
- [SENTRY-1557](#) - getRolesForGroups(),getRoleNamesForGroups() does too many trips to the DB
- [SENTRY-1564](#) - Improve error detection and reporting in MetastoreCacheInitializer.java
- [SENTRY-1577](#) - Support "create function using jar" for hive when Sentry is enabled
- [SENTRY-1581](#) - Provide Log4J metrics reporter
- [SENTRY-1582](#) - Additional comments to clarify the intent of string manipulation methods in SentryStore.java
- [SENTRY-1594](#) - TransactionBlock should become generic
- [SENTRY-1599](#) - CloseablePersistenceManager is no longer needed
- [SENTRY-1609](#) - DelegateSentryStore is subject to JDQL injection
- [SENTRY-1615](#) - SentryStore should not allocate empty objects that are immediately returned
- [SENTRY-1625](#) - PrivilegeOperatePersistence can use QueryParamBuilder Fix [SQL Injection: JDO] in [sentry-provider/sentry-provider-db/src/main/java/org/apache/sentry/provider/db/generic/service/persistent/PrivilegeOperatePersistence.java]
- [SENTRY-1636](#) - Remove thrift dependency on fb303
- [SOLR-9284](#) - The HDFS BlockDirectoryCache should not let it's keysToRelease or names maps grow indefinitely.

- [SOLR-9330](#) - Fix AlreadyClosedException on admin/mbeans?stats=true
- [SOLR-9389](#) - HDFS Transaction logs stay open for writes which leaks Xceivers.
- [SOLR-9506](#) - cache IndexFingerprint for each segment
- [SOLR-9699](#), [SOLR-4668](#) - fix exception from core status in parallel with core reload
- [SOLR-9819](#) - Upgrade Apache commons-fileupload to 1.3.2, fixing a security vulnerability
- [SOLR-9859](#) - backport of replication.properties cannot be updated after being written and neither replication.properties nor index.properties are durable after a crash. Don't log error on NoSuchFileException
- [SOLR-9901](#) - backport of SOLR-9899 Implement move in HdfsDirectoryFactory. SOLR-9899: StandardDirectoryFactory should use optimizations for all FilterDirectorys not just NRTCachingDirectory.
- [SOLR-10031](#) - Validation of filename params in ReplicationHandler
- [SOLR-10114](#), [SOLR-9941](#) - Reordered delete-by-query can delete or omit child documents
- [SOLR-10119](#) - TestReplicationHandler assertion fixes part of
- [SOLR-10121](#), [SOLR-10116](#) - BlockCache corruption with high concurrency
- [SPARK-4105](#) - retry the fetch or stage if shuffle block is corrupt
- [SPARK-12241](#) - [YARN] Improve failure reporting in Yarn client obtainTokenForHBase()
- [SPARK-12523](#) - [YARN] Support long-running of the Spark On HBase and hive meta store.
- [SPARK-16654](#) - [CORE] Add UI coverage for Application Level Blacklisting. Automatically Kill Executors and Nodes when they are Blacklisted
- [SPARK-17874](#) - [CORE] Add SSL port configuration.
- [SPARK-18372](#) - [SQL][BRANCH-1.6] Staging directory fail to be removed
- [SPARK-18750](#) - [YARN] Avoid using "mapValues" when allocating containers.
- [SPARK-19033](#) - [CORE] Add admin acls for history server
- [SPARK-19220](#) - [UI] Make redirection to HTTPS apply to all URIs.
- [SPARK-19554](#) - [UI,YARN] Allow SHS URL to be used for tracking in YARN RM.
- [SPARK-19652](#) - [UI] Do auth checks for REST API access.
- [SPARK-19720](#) - [CORE] Redact sensitive information from SparkSubmit console
- [SQOOP-2349](#) - Add command line option for setting transaction isolation levels for metadata queries
- [SQOOP-2896](#) - Sqoop exec job fails with SQLException Access denied for user
- [SQOOP-2909](#) - Oracle related ImportTest fails after SQOOP-2737
- [SQOOP-2911](#) - Fix failing HCatalogExportTest caused by SQOOP-2863
- [SQOOP-2950](#) - Sqoop trunk has consistent UT failures - need fixing
- [SQOOP-3053](#) - Create a cmd line argument for sqoop.throwOnError and use it through SqoopOptions
- [SQOOP-3055](#) - Fixing MySQL tests failing due to ignored test inputs/configuration
- [SQOOP-3057](#) - Fixing 3rd party Oracle tests failing due to invalid case of column names
- [SQOOP-3066](#) - Introduce an option + env variable to enable/disable SQOOP-2737 feature
- [SQOOP-3068](#) - Enhance error (tool.ImportTool: Encountered IOException running import job: java.io.IOException: Expected schema) to suggest workaround
- [SQOOP-3069](#) - Get OracleExportTest#testUpsertTestExport in line with SQOOP-3066
- [SQOOP-3071](#) - Fix OracleManager to apply localTimeZone correctly in case of Date objects too
- [SQOOP-3072](#) - Reenable escaping in ImportTest#testProductWithWhiteSpaceImport for proper execution
- [SQOOP-3081](#) - use OracleEscapeUtils.escapeIdentifier in OracleUpsertOutputFormat instead of inline appending quotes
- [SQOOP-3124](#) - Fix ordering in column list query of PostgreSQL connector to reflect the logical order instead of adhoc ordering
- [ZOOKEEPER-1045](#) - Backport — Support Quorum Peer mutual authentication via SASL
- [ZOOKEEPER-2044](#) - Backport — Fix CancelledKeyException on ZooKeeper server (zookeeper branch-3.4)

Cloudera Search

Security vulnerability resolved

An attacker can download any file accessible to Solr superuser using the path manipulation vulnerability in Solr (CVE-2017-3163; see [Cloudera Security Bulletins](#) more information). This release provides the fix for it. In addition,

attackers can leverage this vulnerability to manipulate files unrelated to Solr index, by launching a “man in the middle” attack between individual Solr servers. To avoid this issue, enable TLS and Kerberos authentication.

Issues Fixed in CDH 5.10.x

The following topics describe issues fixed in CDH 5.10.x, from newest to oldest release. You can also review [What's New in CDH 5.10.x](#) on page 1824 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.10.2

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.10.2:

- [FLUME-2798](#) - Malformed Syslog messages can lead to OutOfMemoryException
- [FLUME-3080](#) - Close failure in HDFS Sink might cause data loss
- [FLUME-3085](#) - HDFS Sink can skip flushing some BucketWriters, might lead to data loss
- [HADOOP-11400](#) - GraphiteSink does not reconnect to Graphite after 'broken pipe'
- [HADOOP-11599](#) - Client#getTimeout should use IPC_CLIENT_PING_DEFAULT when IPC_CLIENT_PING_KEY is not configured
- [HADOOP-12672](#) - RPC timeout should not override IPC ping interval
- [HADOOP-12751](#) - While using kerberos Hadoop incorrectly assumes names with '@' to be non-simple.
- [HADOOP-13503](#) - Improve SaslRpcClient failure logging
- [HADOOP-13749](#) - KMSClientProvider combined with KeyProviderCache can result in wrong UGI being used
- [HADOOP-13826](#) - S3A Deadlock in multipart copy due to thread pool limits
- [HADOOP-14050](#) - Add process name to kms process
- [HADOOP-14083](#) - Add missing file
hadoop-common-project/hadoop-kms/src/main/tomcat/catalina-default.properties
- [HADOOP-14083](#) - KMS should support old SSL clients.
- [HADOOP-14104](#) - Client should always ask namenode for kms provider path
- [HADOOP-14141](#) - Store KMS SSL keystore password in catalina.properties
- [HADOOP-14195](#) - CredentialProviderFactory\$getProviders is not thread-safe
- [HADOOP-14242](#) - Make KMS Tomcat SSL property sslEnabledProtocols and clientAuth configurable
- [HDFS-10715](#) - NPE when applying AvailableSpaceBlockPlacementPolicy
- [HDFS-11390](#) - Add process name to httpfs process
- [HDFS-11418](#) - HttpFS should support old SSL clients
- [HDFS-11515](#) - -du throws ConcurrentModificationException
- [HDFS-11579](#) - Make HttpFS Tomcat SSL property sslEnabledProtocols and clientAuth configurable
- [HDFS-11689](#) - New exception thrown by DFSClient%isHDFSEncryptionEnabled broke hacky hive code
- [MAPREDUCE-6165](#) - [JDK8] TestCombineFileInputFormat failed on JDK8
- [MAPREDUCE-6201](#) - TestNetworkedJob fails on trunk
- [MAPREDUCE-6839](#) - TestRecovery.testCrashed failed
- [YARN-3251](#) - Fixed a deadlock in CapacityScheduler when computing absoluteMaxAvailableCapacity in LeafQueue
- [YARN-6042](#) - Dump scheduler and queue state information into FairScheduler DEBUG log.
- [YARN-6264](#) - AM not launched when a single vcore is available on the cluster.
- [YARN-6359](#) - TestRM#testApplicationKillAtAcceptedState fails rarely due to race condition
- [YARN-6360](#) - Prevent FS state dump logger from cramming other log files
- [YARN-6453](#) - fairscheduler-statedump.log gets generated regardless of service
- [YARN-6615](#) - AmIpFilter drops query parameters on redirect
- [HBASE-15837](#) - Memstore size accounting is wrong if postBatchMutate() throws exception
- [HBASE-15941](#) - HBase repair should not unsplit healthy splitted region
- [HBASE-16350](#) - Undo server abort from HBASE-14968
- [HBASE-16630](#) - Fragmentation in long running Bucket Cache
- [HBASE-16931](#) - Setting cell's seqId to zero in compaction flow might cause RS down.

- [HBASE-16977](#) - VerifyReplication should log a printable representation of the row keys
- [HBASE-17460](#) - enable_table_replication can not perform cyclic replication of a table
- [HBASE-17501](#) - guard against NPE while reading FileTrailer and HFileBlock
- [HBASE-17574](#) - Clean up how to run tests under hbase-spark module
- [HBASE-17673](#) - Monitored RPC Handler not shown in the WebUI
- [HBASE-17688](#) - MultiRowRangeFilter not working correctly if given same start and stop RowKey
- [HBASE-17710](#) - HBase in standalone mode creates directories with 777 permission
- [HBASE-17717](#) - Explicitly use "sasI" ACL scheme for hbase superuser
- [HBASE-17731](#) - Fractional latency reporting in MultiThreadedAction
- [HBASE-17761](#) - Test TestRemoveRegionMetrics.testMoveRegion fails intermittently because of race condition
- [HBASE-17779](#) - disable_table_replication returns misleading message and does not turn off replication
- [HBASE-17798](#) - RpcServer.Listener.Reader can abort due to CancelledKeyException
- [HBASE-17970](#) - Set yarn.app.mapreduce.am.staging-dir when starting MiniMRCluster
- [HBASE-18096](#) - Limit HFileUtil visibility and add missing annotations
- [HIVE-9481](#) - allow column list specification in INSERT statement
- [HIVE-9567](#) - JSON SerDe not escaping special chars when writing char/varchar data
- [HIVE-10329](#) - Hadoop reflectionutils has issues
- [HIVE-11141](#) - Improve RuleRegExp when the Expression node stack gets huge
- [HIVE-11418](#) - Dropping a database in an encryption zone with CASCADE and trash enabled fails
- [HIVE-11428](#) - Performance: Struct IN() clauses are extremely slow
- [HIVE-11671](#) - Optimize RuleRegExp in DPP codepath
- [HIVE-11842](#) - Improve RuleRegExp by caching some internal data structures
- [HIVE-12179](#) - Add option to not add spark-assembly.jar to Hive classpath
- [HIVE-12768](#) - Thread safety: binary sortable serde decimal deserialization
- [HIVE-13390](#) - Partial backport of HIVE-13390. Backported only httpclient 4.5.2 and httpcore 4.4.4 to fix the Apache Hive SSL vulnerability bug.
- [HIVE-14210](#) - ExecDriver should call jobclient.close() to trigger cleanup
- [HIVE-14380](#) - Queries on tables with remote HDFS paths fail in "encryption" checks.
- [HIVE-14564](#) - Column Pruning generates out of order columns in SelectOperator which cause ArrayIndexOutOfBoundsException.
- [HIVE-14819](#) - FunctionInfo for permanent functions shows TEMPORARY FunctionType
- [HIVE-14943](#) - Partial backport of HIVE-14943 - Base Implementation (of [HIVE-10924](#))
- [HIVE-15282](#) - Different modification times are used when an index is built and when its staleness is checked
- [HIVE-15572](#) - Improve the response time for query canceling when it happens during acquiring locks
- [HIVE-15997](#) - Resource leaks when query is cancelled
- [HIVE-16024](#) - MSCK Repair Requires nonstrict hive.mapred.mode
- [HIVE-16156](#) - FileSinkOperator should delete existing output target when renaming
- [HIVE-16175](#) - Possible race condition in InstanceCache
- [HIVE-16297](#) - Improving hive logging configuration variables
- [HIVE-16394](#) - HoS does not support queue name change in middle of session
- [HIVE-16413](#) - Create table as select does not check ownership of the location
- [HIVE-16459](#) - Forward channelInactive to RpcDispatcher
- [HIVE-16593](#) - SparkClientFactory.stop may prevent JVM from exiting
- [HIVE-16646](#) - Alias in transform ... as clause shouldn't be case sensitive
- [HIVE-16660](#) - Not able to add partition for views in hive when sentry is enabled
- [HIVE-16693](#) - beeline "source" command freezes if you have a comment in it?
- [HUE-5659](#) - [home] Ignore history dependencies when importing document from different cluster
- [HUE-5684](#) - [oozie] Remove duplicates in node data for graph and disable graph tab for large workflows
- [HUE-5714](#) - [hive] Close SQL canary query "Select "Hello World""
- [HUE-5742](#) - [core] Allow user to provide schema name for database via ini

- [HUE-5816](#) - Changing default setting as "allowed_hosts=*"
- [HUE-5873](#) - [editor] The download in progress modal doesn't close automatically on IE11
- [HUE-5984](#) - [search] Escaping corrupts link-meta for building external links in grid dashboard
- [HUE-6045](#) - [pig] Progress bar doesn't turn green on completion
- [HUE-6075](#) - [oozie] Remove email body while displaying external graphs in dashboard
- [HUE-6090](#) - Hue to do a keep alive on idle sessions to HiveServer2
- [HUE-6103](#) - [fb] Log filesystem initialization exceptions
- [HUE-6104](#) - [aws] Check if boto configuration section exists before adding it
- [HUE-6109](#) - [core] Remove the restriction on Document2 invalid chars
- [HUE-6115](#) - [core] Document paths have encoding problems
- [HUE-6131](#) - [hive] Select partition values based on the actual datatypes of the partition column
- [HUE-6133](#) - [home] Typing on the search box crashes IE 11
- [HUE-6144](#) - [oozie] Add generic XSL template to workflow graph parserv
- [HUE-6161](#) - [doc2] Move document conversion upgrade into a migration
- [HUE-6193](#) - [converter] Retain last_executed time when creating doc2 object
- [HUE-6197](#) - [impala] Fix XSS Vulnerability in the old editor error messages
- [HUE-6212](#) - [oozie] Prevent XSS injection in coordinator cron frequency field
- [HUE-6228](#) - [core] Disable touchscreen detection on Nicescroll6228
- [HUE-6250](#) - [frontend] Losing # fragment of full URL on login redirect
- [HUE-6261](#) - [oozie] Avoid JS error preventing workflow action status update
- [HUE-6262](#) - [core] Converter should separate history docs from saved docs
- [HUE-6263](#) - [converter] Delete Doc2 object incase of exception
- [HUE-6264](#) - [converter] Decrease memory usage for users with very high document1 objects
- [HUE-6266](#) - [converter] Remove unnecessary call to document link
- [HUE-6295](#) - [doc2] Avoid unrelated DB calls in sync_documents after import
- [HUE-6310](#) - [doc2] Create missing doc1 links for delete and copy operations
- [HUE-6407](#) - [pig] Play button doesn't come back after killing the running pig job
- [HUE-6446](#) - [oozie] User cant edit shared coordinator or bundle
- [HUE-6604](#) - [oozie] Fix timestamp conversion to server timezone
- [IMPALA-3641](#) - Fix catalogd RPC responses to DROP IF EXISTS.
- [IMPALA-4088](#) - Assign fix values to the minicluster server ports
- [IMPALA-4293](#) - query profile should include error log
- [IMPALA-4544](#) - ASAN should ignore SEGV and leaks
- [IMPALA-4546](#) - Fix Moscow timezone conversion after 2014
- [IMPALA-4615](#) - Fix create_table.sql command order
- [IMPALA-4631](#) - avoid DCHECK in PlanFragementExecutor::Close().
- [IMPALA-4716](#) - Expr rewrite causes IllegalStateException
- [IMPALA-4722](#) - Disable log caching in test_scratch_disk
- [IMPALA-4725](#) - Query option to control Parquet array resolution.
- [IMPALA-4733](#) - Change HBase ports to non-ephemeral
- [IMPALA-4738](#) - STDDEV_SAMP should return NULL for single record input
- [IMPALA-4787](#) - Optimize APPX_MEDIAN() memory usage
- [IMPALA-4822](#) - Implement dynamic log level changes
- [IMPALA-4899](#) - Fix parquet table writer dictionary leak
- [IMPALA-4902](#) - Copy parameters map in HdfsPartition.toThrift().
- [IMPALA-4920](#) - custom cluster tests: fix generation of py.test options
- [IMPALA-4998](#) - Fix missing table lock acquisition.
- [IMPALA-5021](#) - Fix count(*) remaining rows overflow in Parquet.
- [IMPALA-5028](#) - Lock table in /catalog_objects endpoint.
- [IMPALA-5055](#) - Fix DCHECK in parquet-column-readers.cc ReadPageHeader()

- [IMPALA-5088](#) - Fix heap buffer overflow
- [IMPALA-5115](#) - Handle status from HdfsTableSink::WriteClusteredRowBatch
- [IMPALA-5145](#) - Do not constant fold null in CastExprs
- [IMPALA-5154](#) - Handle 'unpartitioned' Kudu tables
- [IMPALA-5156](#) - Drop VLOG level passed into Kudu client
- [IMPALA-5172](#) - Buffer overrun for Snappy decompression
- [IMPALA-5183](#) - increase write wait timeout in BufferedBlockMgrTest
- [IMPALA-5186](#) - Handle failed CreateAndOpenScanner() in MT scan.
- [IMPALA-5189](#) - Pin version of setup-tools-scm
- [IMPALA-5193](#) - Initialize decompressor before finding first tuple
- [IMPALA-5197](#) - Erroneous corrupted Parquet file message
- [IMPALA-5198](#) - Error messages are sometimes dropped before reaching client
- [IMPALA-5208](#) - Bump toolchain to include fixes for and IMPALA-5187
- [IMPALA-5217](#) - KuduTableSink checks null constraints incorrectly
- [IMPALA-5244](#) - test_hdfs_file_open_fail fails on local filesystem build
- [IMPALA-5252](#) - Fix crash in HiveUdfCall::GetStringVal() when mem_limit exceeded
- [IMPALA-5253](#) - Use appropriate transport for StateStoreSubscriber
- [IMPALA-5287](#) - Test skip.header.line.count on gzip
- [IMPALA-5297](#) - Reduce free-pool-test mem requirement to avoid OOM
- [IMPALA-5301](#) - Set Kudu minicluster memory limit
- [IMPALA-5318](#) - Generate access events with fully qualified table names
- [IMPALA-5322](#) - Fix a potential crash in Frontend & Catalog JNI startup
- [IMPALA-5487](#) - Race in runtime-profile.cc::toThrift() can lead to corrupt profiles being generated while query is running
- [IMPALA-5172](#) - fix incorrect cast in call to LZO decompress
- [OOZIE-2739](#) - Remove property expansion pattern from ShellMain's log4j properties content
- [OOZIE-2818](#) - Can't overwrite oozie.action.max.output.data on a per-workflow basis
- [OOZIE-2819](#) - Make Oozie REST API accept multibyte characters for script Actions
- [OOZIE-2844](#) - Increase stability of Oozie actions when log4j.properties is missing or not readable
- [OOZIE-2872](#) - Address backward compatibility issue introduced by OOOZIE-2748
- [OOZIE-2908](#) - Fix typo in oozie.actions.null.args.allowed property in oozie-default.xml
- [SENTRY-1390](#) - Add test cases to ensure usability of URI privileges for HMS binding
- [SENTRY-1422](#) - JDO deadlocks while processing grant while a background thread processes Notificationlogs
- [SENTRY-1512](#) - Refactor the database transaction management Backport Refactor the Sentry database transaction management
- [SENTRY-1476](#) - SentryStore is subject to JDQL injection (Alexander Kolbasov via Vamsee Yarlagadda) Sentry is subject to JDOQL injection
- [SENTRY-1505](#) - CommitContext isn't used by anything and should be removed (CommitContext isn't used by anything and should be removed)
- [SENTRY-1515](#) - Cleanup exception handling in SentryStore
- [SENTRY-1517](#) - SentryStore should actually use function getMSentryRole to get roles
- [SENTRY-1557](#) - getRolesForGroups(), getRoleNamesForGroups() does too many trips to the the DB
- [SENTRY-1594](#) - TransactionBlock should become generic
- [SENTRY-1609](#) - DelegateSentryStore is subject to JDQL injection DelegateSentryStore is subject to JDQL injection
- [SENTRY-1615](#) - SentryStore should not allocate empty objects that are immediately returned
- [SENTRY-1625](#) - PrivilegeOperatePersistence can use QueryParamBuilder Fix [SQL Injection: JDO] in [sentry-provider/sentry-provider-db/src/main/java/org/apache/sentry/provider/db/generic/service/persistent/PrivilegeOperatePersistence.java]
- [SENTRY-1636](#) - Remove thrift dependency on fb303
- [SENTRY-1683](#) - MetastoreCacheInitializer has a race condition in handling results list
- [SENTRY-1714](#) - MetastorePlugin.java should quietly return from renameAuthzObject() when both paths are null
- [SENTRY-1759](#) - UpdatableCache leaks connections

- [SOLR-5776](#) - backportEnabled SSL tests can easily exhaust random generator entropy and block. Set the server side to SHA1PRNG as in Steve's original patch. Use less SSL in a test run. refactor SSLConfig so that SSLTestConfig can provide SSLContexts using a NullSecureRandom to prevent SSL tests from blocking on entropy starved machines Alternate (psuedo random) NullSecureRandom for Constants.SUN_OS replace NullSecureRandom w/ NotSecurePsuedoRandom
- [SOLR-8836](#) - Return 400, and a SolrException when an invalid json is provided to the update handler instead of 500.
- [SOLR-9153](#) - Update Apache commons beanutils version to 1.9.2
- [SOLR-9527](#) - Improve distribution of replicas when restoring a collection
- [SOLR-9836](#) - Add ability to recover from leader when index corruption is detected on SolrCore creation.
- [SOLR-9848](#) - Lower solr.cloud.wait-for-updates-with-stale-state-pause back down from 7 seconds.
- [SOLR-10076](#) - ofHide keystore and truststore passwords from /admin/info/* outputs. - Updated JMXJsonServlet to use new API - Set redaction default true
- [SOLR-10338](#) - backportConfigure SecureRandom non blocking for tests.
- [SOLR-10360](#) - Remove an extra space from Hadoop distcp cmd used by Solr backup/restore
- [SOLR-10430](#) - Add ls command to ZkCLI for listing sub-dirs
- [SPARK-13693](#) - [STREAMING][TESTS] Stop StreamingContext before deleting checkpoint dir
- [SPARK-14930](#) - [SPARK-13693] Fix race condition in CheckpointWriter.stop()
- [SPARK-18922](#) - [SQL][CORE][STREAMING][TESTS] Fix all identified tests failed due to path and resource-not-closed problems on Windows
- [SPARK-19178](#) - [SQL][Backport-to-1.6] convert string of large numbers to int should return null
- [SPARK-19263](#) - DAGScheduler should avoid sending conflicting task set.
- [SPARK-19537](#) - Move pendingPartitions to ShuffleMapStage.
- [SPARK-20922](#) - [CORE][HOTFIX] Don't use Java 8 lambdas in older branches.
- [SPARK-20922](#) - [CORE] Add whitelist of classes that can be deserialized by the launcher.
- [SQOOP-2896](#) - Sqoop exec job fails with SQLException Access denied for user
- [SQOOP-3123](#) - Introduce escaping logic for column mapping parameters (same what Sqoop already uses for the DB column names), thus special column names (e.g. containing '#' character) and mappings related to those columns can be in the same format (thus not confusing the end users), and also eliminates the related AVRO format clashing issues.
- [SQOOP-3140](#) - Removing deprecated mapred.map.max.attempts, mapred.reduce.max.attempts entries and using the new constants directly from Hadoop instead
- [SQOOP-3159](#) - Sqoop (export + --table) with Oracle table_name having '\$' fails with error
- [ZOOKEEPER-2044](#) - Back port to fix CancelledKeyException on ZooKeeper server.

Issues Fixed in CDH 5.10.1

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.10.1:

- [FLUME-2812](#) - Fix semaphore leak causing java.lang.Error: Maximum permit count exceeded in MemoryChannel.
- [FLUME-2889](#) - Fixes to DateTime computations.
- [FLUME-2889](#) - Fixes to DateTime computations.
- [FLUME-2999](#) - Kafka channel and sink should enable statically assigned partition per event via header.
- [FLUME-3027](#) - Change Kafka Channel to clear offsets map after commit.
- [FLUME-3049](#) - Make HDFS sink rotate more reliably in secure mode.
- [HADOOP-11619](#) - FTPFileSystem should override getDefaultPort.
- [HADOOP-12655](#) - TestHttpServer.testBindAddress bind port range is wider than expected.
- [HADOOP-13433](#) - Race in UGI.reloginFromKeytab.
- [HADOOP-13627](#) - Have an explicit KerberosAuthException for UGI to throw, text from public constants.
- [HADOOP-13805](#) - UGI.getCurrentUser() fails if user does not have a keytab associated.
- [HADOOP-13903](#) - Improvements to KMS logging to help debug authorization errors.
- [HADOOP-13911](#) - Remove TRUSTSTORE_PASSWORD related scripts from KMS.

- [HADOOP-13953](#) - Make FTPFileSystem's data connection mode and transfer mode configurable.
- [HADOOP-14003](#) - Make additional KMS tomcat settings configurable.
- [HADOOP-14114](#) - S3A can no longer handle unencoded + in URIs.
- [HDFS-11160](#) - VolumeScanner reports write-in-progress replicas as corrupt incorrectly.
- [HDFS-11275](#) - Check groupEntryIndex and throw a helpful exception on failures when removing ACL.
- [HDFS-11292](#) - log lastWrittenTxId etc info in logSyncAll.
- [HDFS-11306](#) - Print remaining edit logs from buffer if edit log can't be rolled.
- [HDFS-11363](#) - Need more diagnosis info when seeing Slow waitForAckedSeqno.
- [HDFS-11379](#) - DFSInputStream may infinite loop requesting block locations.
- [MAPREDUCE-5155](#) - Race condition in test case TestFetchFailure cause it to fail.
- [MAPREDUCE-6172](#) - TestDbClasses timeouts are too aggressive.
- [MAPREDUCE-6571](#) - JobEndNotification info logs are missing in AM container syslog.
- [MAPREDUCE-6817](#) - The format of job start time in JHS is different from those of submit and finish time.
- [YARN-2306](#) - Add test for leakage of reservation metrics in fair scheduler.
- [YARN-2336](#) - Fair scheduler's REST API returns a missing '[' bracket JSON for deep queue tree.
- [YARN-3269](#) - Yarn.nodemanager.remote-app-log-dir could not be configured to fully qualified path.
- [YARN-3933](#) - FairScheduler: Multiple calls to completedContainer are not safe.
- [YARN-3957](#) - FairScheduler NPE In FairSchedulerQueueInfo causing scheduler page to return 500.
- [YARN-4363](#) - In TestFairScheduler, testcase should not create FairScheduler redundantly.
- [YARN-4544](#) - All the log messages about rolling monitoring interval are shown with WARN level.
- [YARN-4555](#) - TestDefaultContainerExecutor#testContainerLaunchError fails on non-english locale environment.
- [YARN-5136](#) - Error in handling event type APP_ATTEMPT_REMOVED to the scheduler.
- [YARN-5308](#) - FairScheduler: Move continuous scheduling related tests to TestContinuousScheduling.
- [YARN-5752](#) - TestLocalResourcesTrackerImpl#testLocalResourceCache times out.
- [YARN-5859](#) - TestResourceLocalizationService#testParallelDownloadAttemptsForPublicResource sometimes fails.
- [YARN-5890](#) - FairScheduler should log information about AM-resource-usage and max-AM-share for queues.
- [YARN-5920](#) - Fix deadlock in TestRMHA.testTransitionedToStandbyShouldNotHang,
- [YARN-6151](#) - FS preemption does not consider child queues over fairshare if the parent is under.
- [YARN-6175](#) - FairScheduler: Negative vcore for resource needed to preempt.
- [HBASE-12949](#) - Scanner can be stuck in infinite loop if the HFile is corrupted.
- [HBASE-15125](#) - BackportHBaseFsck's adoptHdfsOrphan function creates region with wrong end key boundary.
- [HBASE-15328](#) - sanity check the redirect used to send master info requests to the embedded regionserver.
- [HBASE-15378](#) - Scanner cannot handle heartbeat message with no results.
- [HBASE-15587](#) - FSTableDescriptors.getDescriptor() logs stack trace erroneously.
- [HBASE-15931](#) - Add log for long-running tasks in AsyncProcess HBASE-16289 AsyncProcess stuck messages need to print region/server.
- [HBASE-15955](#) - Disable action in CatalogJanitor#setEnabled should wait for active cleanup scan to finish.
- [HBASE-16032](#) - Possible memory leak in StoreScanner.
- [HBASE-16062](#) - Improper error handling in WAL Reader/Writer creation.
- [HBASE-16237](#) - Blocks for hbase:meta table are not cached in L1 cache.
- [HBASE-16238](#) - It's useless to catch SESSIONEXPIRED exception and retry in RecoverableZooKeeper.
- [HBASE-16266](#) - Do not throw ScannerTimeoutException when catch UnknownScannerException.
- [HBASE-16304](#) - HRegion#RegionScannerImpl#handleFileNotFoundExpection may lead to deadlock when trying to obtain write lock on updatesLock.
- [HBASE-16429](#) - FSHLog: deadlock if rollWriter called when ring buffer filled with appends.
- [HBASE-16460](#) - Can't rebuild the BucketAllocator's data structures when BucketCache uses FileIOEngine.
- [HBASE-16604](#) - Scanner retries on IOException can cause the scans to miss data.
- [HBASE-16649](#) - Truncate table with splits preserved can cause both data loss and truncated data appeared again.
- [HBASE-16662](#) - Fix open POODLE vulnerabilities.
- [HBASE-16721](#) - Concurrency issue in WAL unflushed seqId tracking.

- [HBASE-16807](#) - RegionServer will fail to report new active Hmaster until HMaster/RegionServer failover.
- [HBASE-16841](#) - Data loss in MOB files after cloning a snapshot and deleting that snapshot.
- [HBASE-16960](#) - RegionServer hang when aborting.
- [HBASE-17020](#) - keylen in midkey() dont computed correctly.
- [HBASE-17023](#) - Region left unassigned due to AM and SSH each thinking others would do the assignment work.
- [HBASE-17044](#) - Fix merge failed before creating merged region leaves meta inconsistent.
- [HBASE-17069](#) - RegionServer writes invalid META entries for split daughters in some circumstances.
- [HBASE-17206](#) - FSHLog may roll a new writer successfully with unflushed entries.
- [HBASE-17241](#) - Avoid compacting already compacted mob files with _del files.
- [HBASE-17265](#) - Region left unassigned in master failover when region failed to open.
- [HBASE-17275](#) - Assign timeout may cause region to be unassigned forever.
- [HBASE-17328](#) - Properly dispose of looped replication peers.
- [HBASE-17381](#) - ReplicationSourceWorkerThread can die due to unhandled exceptions.
- [HBASE-17409](#) - Limit jsonp callback name to prevent xss.
- [HBASE-17452](#) - Failed taking snapshot - region Manifest proto-message too large.
- [HBASE-17522](#) - Handle JVM throwing runtime exceptions when we ask for details on heap usage the same as a correctly returned 'undefined'.
- [HBASE-17558](#) - ZK dumping jsp should escape HTML.
- [HBASE-17561](#) - table status page should escape values that may contain arbitrary characters.
- [HBASE-17675](#) - ReplicationEndpoint should choose new sinks if a SaslException occurs.
- [HIVE-7723](#) - Explain plan for complex query with lots of partitions is slow due to in-efficient collection used to find a matching ReadEntity.
- [HIVE-11594](#) - Analyze Table for column names with embedded spaces.
- [HIVE-11849](#) - NPE in HiveHBaseTableShapshotInputFormat in query with just count(*).
- [HIVE-12349](#) - part of the patch which solves IS NULL queries can trigger NPE for timestamp and date columns.
- [HIVE-12465](#) - Hive might produce wrong results when (outer) joins are merged.
- [HIVE-12780](#) - Fix the output of the history command in Beeline HIVE-12789: Fix output twice in the history command of Beeline.
- [HIVE-12976](#) - MetaStoreDirectSql doesn't batch IN lists in all cases.
- [HIVE-13149](#) - Remove some unnecessary HMS connections from HS2.
- [HIVE-13240](#) - GroupByOperator: Drop the hash aggregates when closing operator.
- [HIVE-13864](#) - Beeline ignores the command that follows a semicolon and comment.
- [HIVE-13866](#) - flatten callstack for directSQL errors.
- [HIVE-13895](#) - HoS start-up overhead in yarn-client mode.
- [HIVE-14693](#) - Some partitions will be left out when partition number is the multiple of the option hive.msck.repair.batch.size.
- [HIVE-14764](#) - Enabling "hive.metastore.metrics.enabled" throws OOM in HiveMetastore.
- [HIVE-14820](#) - RPC server for spark inside HS2 is not getting server address properly.
- [HIVE-15338](#) - Wrong result from non-vectorized DATEDIFF with scalar parameter of type DATE/TIMESTAMP.
- [HIVE-15346](#) - "values temp table" should not be an input.
- [HIVE-15359](#) - skip.footer.line.count doesnt work properly for certain situations.
- [HIVE-15410](#) - WebHCat supports get/set table property with its name containing period and hyphen.
- [HIVE-15485](#) - Addendum toInvestigate the DoAs failure in HoS.
- [HIVE-15485](#) - Investigate the DoAs failure in HoS.
- [HIVE-15517](#) - NOT (x <=> y) returns NULL if x or y is NULL.
- [HIVE-15551](#) - memory leak in directsql for mysql+bonecp specific initialization.
- [HIVE-15735](#) - In some cases, view objects inside a view do not have parents.
- [HIVE-15782](#) - query on parquet table returns incorrect result when hive.optimize.index.filter is set to true.
- [HIVE-15872](#) - The PERCENTILE_APPROX UDAF does not work with empty set.
- [HIVE-16019](#) - Query fails when group by/order by on same column with uppercase name.

- [HIVE-16047](#) - Shouldn't try to get KeyProvider unless encryption is enabled.
- [HUE-4969](#) - [core] Rename ini properties for sasl buffer to be standard.
- [HUE-5310](#) - [search] Use Doc2 modal in search_controller.
- [HUE-5408](#) - [oozie] Support old docs while saving shared workflow.
- [HUE-5476](#) - [core] Fix TTL is_idle middleware check.
- [HUE-5482](#) - [home] Handle multiple home/trash directories by merging them into one.
- [HUE-5533](#) - [home] Improve home page load time
- [HUE-5552](#) - [editor] Japanese improvement for new SQL Editor #466
- [HUE-5602](#) - [jb] Add start time filter in jobs page
- [HUE-5604](#) - [core] Update localization for de, es, fr, ja, ko, zh
- [HUE-5670](#) - [doc2] Prevent exception when doc2 object is not linked to doc1
- [HUE-5717](#) - [backend] Some operating system incorrectly detect javascript mime-type as text/x-js instead of application/javascript
- [HUE-5722](#) - [core] Avoid query redaction when string is None
- [HUE-5742](#) - [core] Allow user to provide schema name for database via ini
- [HUE-5756](#) - [doc2] Workaround for improving the query history search time
- [HUE-5758](#) - [oozie] Fix parsing nodes from XML definition
- [HUE-5769](#) - [oozie] Remove mandatory inclusion of Kill row in the workflow dashobard graph
- [HUE-5823](#) - [editor] Cancel running doc search requests when the query has changed
- [HUE-5850](#) - [sentry] Prevent creating roles with empty names
- [HUE-5958](#) - [pig] Fix unicode errors when handling exceptions
- [HUE-5962](#) - [hiveserver2] Update HiveServerClient user object when opening session
- [IMPALA-2605](#) - Omit the sort and mini stress tests.
- [IMPALA-4055](#) - Speed up to_date() with custom implementation.
- [IMPALA-4263](#) - Fix wrong omission of agg/analytic hash exchanges.
- [IMPALA-4282](#) - Remove max length check for type strings.
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog.
- [IMPALA-4675](#) - Case-insensitive matching of Parquet fields.
- [IMPALA-4702](#) - Fix command line help for webserver_private_key_file.
- [IMPALA-4705](#), [IMPALA-4779](#), [IMPALA-4780](#) - Fix some Expr bugs with codegen.
- [IMPALA-4742](#) - Change "{}".format() to "{0}".format() for Py 2.6.
- [IMPALA-4749](#) - hit DCHECK in sorter with scratch limit.
- [IMPALA-4767](#) - Workaround for HIVE-15653 to preserve table stats.
- [IMPALA-4808](#) - old hash join can reference invalid memory.
- [IMPALA-4828](#) - Alter Kudu schema outside Impala may crash on read.
- [IMPALA-4854](#) - Fix incremental stats with complex types.
- [IMPALA-4916](#) - Fix maintenance of set of item sets in DisjointSet.
- [IMPALA-4981](#) - Re-enable spilling with MT_DOP.
- [IMPALA-4995](#) - Fix integer overflow in TopNNode::PrepareForOutput.
- [IMPALA-4997](#) - Fix overflows in Sorter::TupleIterator.
- [OOZIE-2243](#) - Kill Command does not kill the child job for java action.
- [OOZIE-2519](#) - Oozie HA with SSL info is slightly incorrect.
- [OOZIE-2584](#) - Eliminate Thread.sleep() calls in TestMemoryLocks.
- [OOZIE-2742](#) - Unable to kill applications based on tag.
- [OOZIE-2748](#) - NPE in LauncherMapper.printArgs().
- [OOZIE-2757](#) - Malformed xml is Spark action doc page.
- [OOZIE-2777](#) - Config-default.xml longer than 64k results in java.io.UTFDataFormatException.
- [OOZIE-2787](#) - Oozie distributes application jar twice making the spark job fail.
- [OOZIE-2802](#) - Spark action failure on Spark 2.1.0 due to duplicate sharelibs.
- [SENTRY-1508](#) - MetastorePlugin.java does not handle properly initialization failure.

- [SENTRY-1520](#) - Provide mechanism for triggering HMS full snapshot.
- [SENTRY-1564](#) - Improve error detection and reporting in MetastoreCacheInitializer.java.
- [SENTRY-1605](#) - 'SENTRY-1508 need to be fixed because of Kerberos initialization issue.
- [SOLR-9284](#) - The HDFS BlockDirectoryCache should not let its keysToRelease or names maps grow indefinitely.
- [SOLR-9330](#) - Fix AlreadyClosedException on admin/mbeans?stats=true.
- [SOLR-9699](#), [SOLR-4668](#) - fix exception from core status in parallel with core reload.
- [SOLR-9819](#) - Upgrade Apache commons-fileupload to 1.3.2, fixing a security vulnerability.
- [SOLR-9859](#) - backport of replication.properties cannot be updated after being written and neither replication.properties or index.properties are durable in the face of a crash. Don't log error on NoSuchFileException.
- [SOLR-9901](#) - backport of, SOLR-9899Implement move in HdfsDirectoryFactory. SOLR-9899: StandardDirectoryFactory should use optimizations for all FilterDirectorys not just NRTCachingDirectory.
- [SOLR-10031](#) - Validation of filename params in ReplicationHandler.
- [SOLR-10114](#), [SOLR-9941](#) - Reordered delete-by-query can delete or omit child documents
- [SOLR-10119](#) - TestReplicationHandler assertion fixes part of.
- [SOLR-10121](#), [SOLR-10116](#) - BlockCache corruption with high concurrency.
- [SPARK-12241](#) - [YARN] Improve failure reporting in Yarn client obtainTokenForHBase().
- [SPARK-12523](#) - [YARN] Support long-running of the Spark On HBase and hive meta store.
- [SPARK-18750](#) - [YARN] Follow up: move test to correct directory in 2.1 branch.
- [SPARK-18750](#) - [YARN] Avoid using "mapValues" when allocating containers.
- [SQOOP-2349](#) - Add command line option for setting transaction isolation levels for metadata queries.
- [SQOOP-2896](#) - Sqoop exec job fails with SQLException Access denied for user.
- [SQOOP-2909](#) - Oracle related ImportTest fails after SQOOP-2737.
- [SQOOP-2911](#) - Fix failing HCatalogExportTest caused by SQOOP-2863.
- [SQOOP-2950](#) - Sqoop trunk has consistent UT failures - need fixing.
- [SQOOP-3053](#) - Create a cmd line argument for sqoop.throwOnError and use it through SqoopOptions.
- [SQOOP-3055](#) - Fixing MySQL tests failing due to ignored test inputs/configuration.
- [SQOOP-3057](#) - Fixing 3rd party Oracle tests failing due to invalid case of column names.
- [SQOOP-3068](#) - Enhance error (tool.ImportTool: Encountered IOException running import job: java.io.IOException: Expected schema) to suggest workaround.
- [SQOOP-3071](#) - Fix OracleManager to apply localTimeZone correctly in case of Date objects too
- [SQOOP-3072](#) - Reenable escaping in ImportTest#testProductWithWhiteSpaceImport for proper execution.
- [SQOOP-3081](#) - use OracleEscapeUtils.escapeldentifier in OracleUpsertOutputFormat instead of inline appending quotes.
- [SQOOP-3124](#) - Fix ordering in column list query of PostgreSQL connector to reflect the logical order instead of adhoc ordering.

Issues Fixed in CDH 5.10.0

CDH 5.10.0 fixes the following issues.

Apache Hadoop Common

Have Hadoop use commons-daemon version from root pom

Apache Commons Daemon 1.0.3 has been removed in favor of Apache Commons Daemon 1.0.13.

Apache HBase

HBase shell commands correctly return missing namespace name

In prior releases, passing a non-existent namespace with GRANT and REVOKE commands to the HBase shell returned an error message displaying the user name, rather than the name of the namespace. In this release, the error message correctly displays the missing namespace name. This change is a result of an underlying change to the error handling for the HBase shell; specifically, the boolean `namespace_exists?` method now returns `false` when the namespace is not found, rather than the `NamespaceNotFoundException`.

Apache HDFS

*fsck should also report decommissioning replicas*Bug: [HDFS-7933](#)

Cloudera Bug: CDH-47416

The output of HDFS fsck now also contains information about decommissioning replicas.

*du reports false used space after appending to snapshoted files and deleting them*Bug: [HDFS-7933](#)

Cloudera Bug: CDH-40502

Disk usage summaries previously incorrectly counted files twice if they had been renamed (including files moved to Trash) since the last snapshot. Summaries now include current data plus snapshot data that is no longer under the directory either due to deletion or being moved outside of the directory.

Fixed in Versions: 5.10, 5.10, 5.11.0

*ACL inheritance conflicts with umaskmode*Bug: [HDFS-6962](#)

Cloudera Bug: CDH-26136

Previously, HDFS ACLs applied the client `umask` to the permissions when inheriting default ACLs from the parent directory. This differs from the POSIX ACL specification. HDFS can now ignore the client `umask` to comply with the POSIX ACL specification.

Because this is a [backward-incompatible](#) change, this behavior is disabled by default. To enable it, set `dfs.namenode.posix.acl.inheritance.enabled` to `true` in `hdfs-site.xml`.

Hue

Security and Administration

- [HUE-4541](#), CDH-44287 - Support Kerberos mutual authentication across HUE
- [HUE-4969](#), CDH-34544 - Support HiveServer2 + SASL (`hive.server2.thrift.sasl.qop="auth-conf"`)
- [HUE-4372](#), CDH-42194 - Turn off HSTS header in Hue Load Balancer (server only)
- [HUE-3079](#), CDH-41093 - Display YARN jobs and logs from JHS when not in RM
- **Cloudera Bug:** OPSAPS-34608 - Support Oracle DB connectivity for CDH package installations

SQL Editor and General UX

- [HUE-5070](#), CDH-40886 - Stabilize Import documents
- [HUE-4032](#), CDH-44094 - Improve Sample Popup performance
- [HUE-4073](#) - Optimize SQL Assist rendering for large scrolls
- [HUE-4039](#) - Refine Autocompleter results
- [HUE-4726](#), CDH-35007 - Improve viewing Parquet formatted data

Apache Oozie

*After enabling HA, Oozie may fail to start due to "NoSuchFieldError: EXTERNAL_PROPERTY"*Bug: [HIVE-1640](#)

Cloudera Bug: CDH-42408

This issue happens in rare cases. Due to an incompatibility with the version of Jackson used by Oozie and Hive, and depending on the order that jars are loaded into Oozie's classpath, Oozie may fail to start.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.10.0:

- [AVRO-1584](#) - Java: Escape characters not allowed in JSON in toString

- [AVRO-1642](#) - Java: Do not generate invalid all-args constructor
- [AVRO-1684](#) - Add time types to the specific compiler
- [AVRO-1799](#) - Fix GenericRecord#toString ByteBuffer bug
- [AVRO-1847](#) - IDL compiler should use BigDecimal to represent decimal logical type
- [AVRO-1869](#) - Java: Fix Decimal conversion from ByteBuffer
- [AVRO-1877](#) - Restore correct javaUnbox in Specific Compiler
- [AVRO-1895](#) - Fix deepCopy() to get correct logical type conversion
- [AVRO-1895](#) - Java: Fix GenericData#deepCopy() to support logical types
- [CRUNCH-616](#) - Replace (possibly copyrighted) Maugham text with Dickens
- [FLUME-2171](#) - Add Interceptor to remove headers from event
- [FLUME-3025](#) - Expose FileChannel.open on JMX
- [FLUME-3031](#) - Change sequence source to reset its counter for event body on channel exception
- [HADOOP-3733](#) - "s3x:" URLs break when Secret Key contains a slash, even if encoded
- [HADOOP-7930](#) - Kerberos relogin interval in UserGroupInformation should be configurable
- [HADOOP-10062](#) - Race condition in MetricsSystemImpl#publishMetricsNow that causes incorrect results
- [HADOOP-11381](#) - Fix findbugs warnings in hadoop-distcp, hadoop-aws, hadoop-azure, and hadoop-openstack
- [HADOOP-11412](#) - POMs mention "The Apache Software License" rather than "Apache License"
- [HADOOP-11520](#) - Clean incomplete multi-part uploads in S3A tests
- [HADOOP-11720](#) - [JDK8] Fix javadoc errors caused by incorrect or illegal tags in hadoop-tools
- [HADOOP-11780](#) - Prevent IPC reader thread death
- [HADOOP-11922](#) - Misspelling of threshold in log4j.properties for tests in hadoop-tools
- [HADOOP-12169](#) - ListStatus on empty dir in S3A lists itself instead of returning an empty list
- [HADOOP-12169](#) - ListStatus on empty dir in S3A lists itself instead of returning an empty list
- [HADOOP-12292](#) - Make use of DeleteObjects optional
- [HADOOP-12325](#) - RPC Metrics : Add the ability track and log slow RPCs
- [HADOOP-12444](#) - Support lazy seek in S3AInputStream
- [HADOOP-12597](#) - In kms-site.xml configuration hadoop.security.keystore.JavaKeyStoreProvider.password should be updated with new name
- [HADOOP-12696](#) - Add tests for S3FileSystem Contract
- [HADOOP-12801](#) - Suppress obsolete S3FileSystem tests
- [HADOOP-12807](#) - S3AFileSystem should read AWS credentials from environment variables
- [HADOOP-12846](#) - Credential Provider Recursive Dependencies
- [HADOOP-12851](#) - S3AFileSystem Uptake of ProviderUtils.excludeIncompatibleCredentialProviders
- [HADOOP-12891](#) - S3AFileSystem should configure Multipart Copy threshold and chunk size
- [HADOOP-12994](#) - Specify PositionedReadable, add contract tests, fix problems
- [HADOOP-13028](#) - Add low level counter metrics for S3A; use in read performance tests
- [HADOOP-13056](#) - Print expected values when rejecting a server's determined principal
- [HADOOP-13065](#) - Add a new interface for retrieving FS and FC Statistics
- [HADOOP-13113](#) - Enable parallel test execution for hadoop-aws
- [HADOOP-13116](#) - Jets3tNativeS3FileSystemContractTest does not run
- [HADOOP-13122](#) - Customize User-Agent header sent in HTTP requests by S3A
- [HADOOP-13130](#) - s3a failures can surface as RTEs, not IOEs
- [HADOOP-13131](#) - Add tests to verify that S3A supports SSE-S3 encryption
- [HADOOP-13139](#) - Branch-2: S3a to use thread pool that blocks clients
- [HADOOP-13145](#) - In DistCp, prevent unnecessary getFileStatus call when not preserving metadata
- [HADOOP-13158](#) - S3AFileSystem#toString might throw NullPointerException due to null cannedACL
- [HADOOP-13162](#) - Consider reducing number of getFileStatus calls in S3AFileSystem.mkdirs
- [HADOOP-13171](#) - Add StorageStatistics to S3A; instrument some more operations
- [HADOOP-13183](#) - S3A proxy tests fail after httpClient/httpcore upgrade
- [HADOOP-13188](#) - S3A file-create should throw error rather than overwrite directories

- [HADOOP-13201](#) - Print the directory paths when ViewFs denies the rename operation on internal dirs
- [HADOOP-13203](#) - S3A: Support fadvise "random" mode for high performance readPositioned() reads
- [HADOOP-13208](#) - S3A listFiles(recursive=true) to do a bulk listObjects instead of walking the pseudo-tree of directories
- [HADOOP-13212](#) - Provide an option to set the socket buffers in S3AFileSystem
- [HADOOP-13237](#) - s3a initialization against public bucket fails if caller lacks any credentials
- [HADOOP-13239](#) - Deprecate s3:// in branch-2
- [HADOOP-13241](#) - Document s3a better
- [HADOOP-13252](#) - Tune S3A provider plugin mechanism
- [HADOOP-13280](#) - FileSystemStorageStatistics#getLong("readOps") should return readOps + largeReadOps
- [HADOOP-13283](#) - Support reset operation for new global storage statistics and per FS storage stats
- [HADOOP-13284](#) - FileSystemStorageStatistics must not attempt to read non-existent rack-aware read stats in branch-2.8
- [HADOOP-13287](#) - TestS3ACredentials#testInstantiateFromURL fails if AWS secret key contains +
- [HADOOP-13288](#) - Guard null stats key in FileSystemStorageStatistics
- [HADOOP-13291](#) - Probing stats in DFSOpsCountStatistics/S3AStorageStatistics should be correctly implemented
- [HADOOP-13305](#) - Define common statistics names across schemes
- [HADOOP-13324](#) - s3a tests don't authenticate with S3 frankfurt (or other V4 auth only endpoints)
- [HADOOP-13368](#) - DFSOpsCountStatistics\$OpType#fromSymbol and s3a.Statistic#fromSymbol should be O(1) operation
- [HADOOP-13387](#) - Users always get told off for using S3 —even when not using it
- [HADOOP-13389](#) - TestS3ATemporaryCredentials.testSTS error when using IAM credentials
- [HADOOP-13396](#) - Allow pluggable audit loggers in KMS
- [HADOOP-13405](#) - Doc for fs.s3a.acl.default indicates incorrect values
- [HADOOP-13406](#) - S3AFileSystem: Consider reusing filestatus in delete() and mkdirs()
- [HADOOP-13442](#) - Optimize UGI group lookups
- [HADOOP-13447](#) - Refactor S3AFileSystem to support introduction of separate metadata repository and tests
- [HADOOP-13512](#) - ReloadingX509TrustManager should keep reloading in case of exception.
- [HADOOP-13541](#) - Explicitly declare the Joda time version S3A depends on
- [HADOOP-13590](#) - Retry until TGT expires even if the UGI renewal thread encountered exception
- [HADOOP-13601](#) - Fix a log message typo in AbstractDelegationTokenSecretManager
- [HADOOP-13641](#) - Update UGI#spawnAutoRenewalThreadForUserCreds to reduce indentation
- [HADOOP-13684](#) - Snappy may complain Hadoop is built without snappy if libhadoop is not found
- [HADOOP-13698](#) - Document caveat for KeyShell when underlying KeyProvider does not delete a key
- [HADOOP-13838](#) - KMSTokenRenewer should close providers
- [HADOOP-13864](#) - KMS should not require truststore password
- [HDFS-742](#) - A down DataNode makes Balancer to hang on repeatedly asking NameNode its partial block list
- [HDFS-2390](#) - dfsadmin -setBalancerBandwidth does not validate -ve value
- [HDFS-4396](#) - Add START_MSG/SHUTDOWN_MSG for ZKFC
- [HDFS-6565](#) - Use jackson instead jetty json in hdfs-client
- [HDFS-7224](#) - Allow reuse of NN connections via webhdfs
- [HDFS-7384](#) - getfacl command and getAclStatus output should be in sync
- [HDFS-7411](#) - Change decommission logic to throttle by blocks rather than nodes in each interval
- [HDFS-7537](#) - Add "UNDER MIN REPL'D BLOCKS" count to fsck
- [HDFS-7933](#) - fsck should also report decommissioning replicas
- [HDFS-8037](#) - CheckAccess in WebHDFS silently accepts malformed FsActions parameters
- [HDFS-8039](#) - Fix TestDebugAdmin#testRecoverLease and testVerfiyBlockChecksumCommand on Windows
- [HDFS-8405](#) - Fix a typo in NamenodeFsck
- [HDFS-8542](#) - WebHDFS getHomeDirectory behavior does not match specification
- [HDFS-8721](#) - Add a metric for number of encryption zones

- [HDFS-8826](#) - In Balancer, add an option to specify the source node list so that balancer only selects blocks to move from those nodes
- [HDFS-8923](#) - Add -source flag to balancer usage message
- [HDFS-8986](#) - Add option to -du to calculate directory space usage excluding snapshots
- [HDFS-9005](#) - Provide support for upgrade domain script
- [HDFS-9019](#) - Adding informative message to sticky bit permission denied exception
- [HDFS-9063](#) - Correctly handle snapshot path for getContentSummary
- [HDFS-9214](#) - Add missing license header
- [HDFS-9214](#) - Support reconfiguring dfs.datanode.balance.max.concurrent.moves without DN restart
- [HDFS-9223](#) - Code cleanup for DatanodeDescriptor and HeartbeatManager
- [HDFS-9257](#) - Improve error message for "Absolute path required" in INode.java to contain the rejected path
- [HDFS-9279](#) - Decommissioned capacity should not be considered for configured/used capacity
- [HDFS-9389](#) - Add maintenance states to AdminStates
- [HDFS-9392](#) - Admins support for maintenance state
- [HDFS-9444](#) - Add utility to find set of available ephemeral ports to ServerSocketUtil
- [HDFS-9500](#) - Fix software version counts for DataNodes during rolling upgrade
- [HDFS-9724](#) - Degraded performance in WebHDFS listing as it does not reuse ObjectMapper
- [HDFS-9745](#) - TestSecureNNWithQJM#testSecureMode sometimes fails with timeouts
- [HDFS-9790](#) - HDFS Balancer should exit with a proper message if upgrade is not finalized
- [HDFS-9839](#) - Reduce verbosity of processReport logging
- [HDFS-9885](#) - Correct the distcp counters name while displaying counters
- [HDFS-9926](#) - MiniDFSCluster leaks dependency Mockito via DataNodeTestUtils
- [HDFS-9934](#) - ReverseXML oiv processor should bail out if the XML file's layoutVersion doesn't match oiv's
- [HDFS-9951](#) - Use string constants for XML tags in OfflinelImageReconstructor
- [HDFS-10276](#) - HDFS should not expose path info that user has no permission to see
- [HDFS-10277](#) - PositionedReadable test testReadFullyZeroByteFile failing in HDFS
- [HDFS-10291](#) - TestShortCircuitLocalRead failing
- [HDFS-10415](#) - TestDistributedFileSystem#MyDistributedFileSystem attempts to set up statistics before initialize() is called
- [HDFS-10423](#) - Increase default value of httpfs maxHTTPHeaderSize
- [HDFS-10505](#) - OIV's ReverseXML processor should support ACLs
- [HDFS-10549](#) - Correctly revoke file leases when closing files
- [HDFS-10553](#) - DiskBalancer: Rename Tools/DiskBalancer class to Tools/DiskBalancerCLI
- [HDFS-10599](#) - DiskBalancer: Execute CLI via Shell
- [HDFS-10627](#) - Volume Scanner marks a block as "suspect" even if the exception is network-related
- [HDFS-10628](#) - Log HDFS Balancer exit message to its own log
- [HDFS-10655](#) - Fix path related byte array conversion bugs
- [HDFS-10656](#) - Optimize conversion of byte arrays back to path string
- [HDFS-10674](#) - Optimize creating a full path from an inode
- [HDFS-10694](#) - processReport() should print blockReportId in each log message
- [HDFS-10738](#) - MR1Fix TestRefreshUserMappings.testRefreshSuperUserGroupsConfiguration test failure
- [HDFS-10747](#) - o.a.h.hdfs.tools.DebugAdmin usage message is misleading
- [HDFS-10756](#) - Expose getTrashRoot to HTTPFS and WebHDFS
- [HDFS-10763](#) - Open files can leak permanently due to inconsistent lease update
- [HDFS-10784](#) - Implement WebHdfsFileSystem#listStatusIterator
- [HDFS-10797](#) - Disk usage summary of snapshots causes renamed blocks to get counted twice
- [HDFS-10807](#) - Doc about upgrading to a version of HDFS with snapshots may be confusing
- [HDFS-10809](#) - getNumEncryptionZones causes NPE in branch-2.7
- [HDFS-10823](#) - Implement HttpFSFileSystem#listStatusIterator
- [HDFS-10832](#) - Propagate ACL bit and isEncrypted bit in HttpFS FileStatus permissions

- [HDFS-10837](#) - Standardize serialization of WebHDFS DirectoryListing
- [HDFS-10870](#) - Wrong dfs.namenode.acls.enabled default in HdfsPermissionsGuide.appt.vm
- [HDFS-10875](#) - Optimize du -x to cache intermediate result
- [HDFS-10876](#) - Dispatcher#dispatch should log IOException stacktrace
- [HDFS-10878](#) - TestDFSClientRetries#testIdempotentAllocateBlockAndClose throws ConcurrentModificationException
- [HDFS-10883](#) - `getTrashRoot`'s behavior is not consistent in DFS after enabling EZ
- [HDFS-10915](#) - Fix time measurement bug in TestDatanodeRestart
- [HDFS-10918](#) - Add a tool to get FileEncryptionInfo from CLI
- [HDFS-10960](#) - TestDataNodeHotSwapVolumes#testRemoveVolumeBeingWritten fails at disk error verification after volume remove
- [HDFS-11009](#) - Add a tool to reconstruct block meta file from CLI
- [HDFS-11015](#) - Enforce timeout in balancer
- [HDFS-11053](#) - Unnecessary superuser check in versionRequest()
- [HDFS-11069](#) - Tighten the authorization of datanode RPC
- [HDFS-11080](#) - Update HttpFS to use ConfigRedactor
- [HDFS-11120](#) - TestEncryptionZones should waitActive
- [HDFS-11229](#) - HDFS-11056 failed to close meta file
- [MAPREDUCE-6497](#) - Fix wrong value of JOB_FINISHED event in JobHistoryEventHandler
- [MAPREDUCE-6541](#) - Exclude scheduled reducer memory when calculating available mapper slots from headroom to avoid deadlock
- [MAPREDUCE-6579](#) - JobStatus#getFailureInfo should not output diagnostic information when the job is running
- [MAPREDUCE-6740](#) - Enforce mapreduce.task.timeout to be at least mapreduce.task.progress-report.interval
- [MAPREDUCE-6750](#) - Fix TestHSAdminServer#testRefreshSuperUserGroups
- [MAPREDUCE-6763](#) - Shuffle server listen queue is too small
- [MAPREDUCE-6764](#) - Teragen LOG initialization bug
- [MAPREDUCE-6765](#) - MR should not schedule container requests in cases where reducer or mapper containers demand resource larger than the maximum supported
- [MAPREDUCE-6776](#) - yarn.app.mapreduce.client.job.max-retries should have a more useful default
- [MAPREDUCE-6789](#) - Fix TestAMWebApp failure
- [YARN-2246](#) - Made the proxy tracking URL always be http(s)://proxy addr:port/proxy/<appld> to avoid duplicate sections
- [YARN-2913](#) - DOCS. Fair scheduler should have ability to set MaxResourceDefault for each queue
- [YARN-2980](#) - Move health check script related functionality to hadoop-common
- [YARN-3094](#) - Reset timer for liveness monitors after RM recovery
- [YARN-3223](#) - Resource update during NM graceful decommission
- [YARN-3239](#) - WebAppProxy does not support a final tracking url which has query fragments and params
- [YARN-3375](#) - NodeHealthScriptRunner.shouldRun() check is performing 3 times for starting NodeHealthScriptRunner
- [YARN-3412](#) - RM tests should use MockRM where possible
- [YARN-3582](#) - NPE in WebAppProxyServlet
- [YARN-3893](#) - Both RM in active state when Admin#transitionToActive failure from refreshAll()
- [YARN-4115](#) - Reduce loglevel of ContainerManagementProtocolProxy to Debug
- [YARN-4132](#) - Separate configs for nodemanager to resourcemanager connection timeout and retries
- [YARN-4201](#) - AMBlacklist does not work for minicluster
- [YARN-4710](#) - Reduce logging application reserved debug info in FSAppAttempt#assignContainer
- [YARN-4743](#) - FairSharePolicy breaks TimSort assumption
- [YARN-4767](#) - Network issues can cause persistent RM UI outage
- [YARN-4794](#) - Deadlock in NMClientImpl
- [YARN-4911](#) - Bad placement policy in FairScheduler causes the RM to crash
- [YARN-4927](#) - TestRMHA#testTransitionedToActiveRefreshFail fails with FairScheduler
- [YARN-5009](#) - NMLevelDbStateStoreService database can grow substantially leading to longer recovery times

- [YARN-5082](#) - Limit ContainerId increase in fair scheduler if the num of node app reserved reached the limit
- [YARN-5197](#) - RM leaks containers if running container disappears from node update
- [YARN-5353](#) - ResourceManager can leak delegation tokens when they are shared across apps
- [YARN-5453](#) - FairScheduler#update may skip update demand resource of child queue/app if current demand reached maxResource
- [YARN-5462](#) - TestNodeStatusUpdater.testNodeStatusUpdaterRetryAndNMShutdown fails intermittently
- [YARN-5616](#) - Clean up WeightAdjuster
- [YARN-5672](#) - FairScheduler: Wrong queue name in log when adding application
- [YARN-5677](#) - RM should transition to standby when connection is lost for an extended period
- [YARN-5693](#) - Reduce loglevel to Debug in ContainerManagementProtocolProxy and AMRMClientImpl
- [YARN-5694](#) - ZKRMStateStore can prevent the transition to standby if the ZK node is unreachable
- [YARN-5736](#) - Addendum. Fixes segfault due to unterminated string
- [YARN-5736](#) - YARN container executor config does not handle white space
- [YARN-5754](#) - Null check missing for earliest in FifoPolicy
- [YARN-5834](#) - TestNodeStatusUpdater.testNMRMConnectionConf compares nodemanager wait time to the incorrect value.
- [YARN-5942](#) - "Overridden" is misspelled as "overriden" in FairScheduler.md
- [HBASE-12940](#) - Expose listPeerConfigs and getPeerConfig to the HBase shell
- [HBASE-15297](#) - Correct handling of namespace existence checks in shell
- [HBASE-15393](#) - Enable table replication command will fail when parent znode is not default in peer cluster
- [HBASE-15526](#) - Make SnapshotManager accessible through MasterServices
- [HBASE-15633](#) - Backport HBASE-15507 to branch-1
- [HBASE-15769](#) - Perform validation on cluster key for add_peer
- [HBASE-16146](#) - Counters are expensive
- [HBASE-16464](#) - archive folder grows bigger and bigger due to corrupt snapshot under tmp dir
- [HBASE-16490](#) - Fix race condition between SnapshotManager and SnapshotCleaner
- [HBASE-16653](#) - Backport HBASE-11393 to branches which support namespace
- [HBASE-17058](#) - Lower epsilon used for jitter verification from HBASE-15324
- [HBASE-17072](#) - CPU usage starts to climb up to 90-100% when using G1GC; purge ThreadLocal usage
- [HIVE-4924](#) - JDBC: Support query timeout for jdbc
- [HIVE-9423](#) - HiveServer2: Provide the user with different error messages depending on the Thrift client exception code
- [HIVE-9518](#) - Implement MONTHS_BETWEEN aligned with Oracle one
- [HIVE-9518](#) - Implement MONTHS_BETWEEN aligned with Oracle one
- [HIVE-9664](#) - Hive 'add jar' command should be able to download and add jars from a repository
- [HIVE-10267](#) - HIVE-9664 makes hive depend on ivysettings.xml : trivial breakage fix
- [HIVE-10276](#) - Implement date_format(timestamp, fmt) UDF
- [HIVE-10276](#) - Implement date_format(timestamp, fmt) UDF
- [HIVE-10576](#) - Add jar command does not work with Windows OS
- [HIVE-10644](#) - Create SHA2 UDF
- [HIVE-10644](#) - Create SHA2 UDF
- [HIVE-11032](#) - Enable more tests for grouping by skewed data
- [HIVE-11538](#) - Add an option to skip init script while running tests
- [HIVE-11920](#) - ADD JAR failing with URL schemes other than file/ivy/hdfs
- [HIVE-12619](#) - Switching the field order within an array of structs causes the query to fail
- [HIVE-12646](#) - Revert "beeline and HIVE CLI do not parse ; in quote properly"
- [HIVE-12653](#) - The property "serialization.encoding" in the class "org.apache.hadoop.hive.contrib.serde2.MultiDelimitSerDe" does not work
- [HIVE-12908](#) - Improve dynamic partition loading III
- [HIVE-12988](#) - Improve dynamic partition loading IV

- [HIVE-13033](#) - SPDO unnecessarily duplicates columns in key & value of mapper output
- [HIVE-13129](#) - CliService leaks HMS connection
- [HIVE-13539](#) - HiveHFileOutputFormat searching the wrong directory for HFiles
- [HIVE-13705](#) - Insert into table removes existing data
- [HIVE-13716](#) - Improve dynamic partition loading V
- [HIVE-13760](#) - Add a HIVE_QUERY_TIMEOUT configuration to kill a query if a query is running for more than the configured timeout value.
- [HIVE-13786](#) - Fix the unit test failure
org.apache.hive.service.cli.session.TestHiveSessionImpl.testLeakOperationHandle
- [HIVE-13904](#) - Ignore case when retrieving ColumnInfo from RowResolver
- [HIVE-13911](#) - Load inpath fails throwing org.apache.hadoop.security.AccessControlException
- [HIVE-13933](#) - Add an option to turn off parallel file moves
- [HIVE-13936](#) - Add streaming support for row_number
- [HIVE-13960](#) - Session may timeout before idle timeout time for synchronous operations
- [HIVE-14011](#) - MessageFactory is not pluggable
- [HIVE-14100](#) - Adding a new logged_in_user() UDF which returns the user provided when connecting
- [HIVE-14175](#) - Fix creating buckets without scheme information
- [HIVE-14301](#) - insert overwrite fails for nonpartitioned tables in s3
- [HIVE-14358](#) - Add metrics for number of queries executed for each execution engine
- [HIVE-14373](#) - Add integration tests for hive on S3
- [HIVE-14444](#) - Upgrade QTest execution framework to junit4 - migrate most of them
- [HIVE-14753](#) - Track the number of open/closed/abandoned sessions in HS2
- [HIVE-14775](#) - Cleanup IOException usage in Metrics APIs
- [HIVE-14784](#) - Operation logs are disabled automatically if the parent directory does not exist
- [HIVE-14822](#) - Add support for credential provider for jobs launched from Hiveserver2
- [HIVE-14924](#) - MSCK REPAIR table with single threaded is throwing null pointer exception
- [HIVE-15000](#) - Remove addlocaldriverjar, and addlocaldrivername from command line help
- [HIVE-15022](#) - Missing hs2-connection-timed-out in BeeLine.properties
- [HIVE-15114](#) - Added few more test cases to validate issue
- [HIVE-15114](#) - Remove extra MoveTask operators from the ConditionalTask
- [HIVE-15121](#) - Last MR job in Hive should be able to write to a different scratch directory
- [HIVE-15199](#) - INSERT INTO data on S3 is replacing the old rows with the new ones
- [HIVE-15226](#) - Add a different masking comment to qtests blobstore output
- [HIVE-15246](#) - Add a making comment to blobstore staging paths on qtest output
- [HIVE-15266](#) - Edit test output of negative blobstore tests to match HIVE-15226
- [HIVE-15280](#) - Hive.mvFile() misses the "." char when joining the filename + extension
- [HIVE-15291](#) - Comparison of timestamp fails if only date part is provided
- [HIVE-15355](#) - Concurrency issues during parallel moveFile due to HDFSUtils.setFullFileStatus
- [HIVE-15361](#) - INSERT dynamic partition on S3 fails with a MoveTask failure
- [HIVE-15363](#) - Execute hive-blobstore tests using ProxyLocalFileSystem
- [HIVE-15367](#) - CTAS with LOCATION should write temp data under location directory rather than database location
- [HIVE-15385](#) - Failure to inherit permissions when running HdfsUtils.setFullFileStatus(..., false) causes queries to fail
- [IMPALA-1169](#) - Admission control info on the queries debug webpage
- [IMPALA-1286](#) - Extract common conjuncts from disjunctions.
- [IMPALA-1430](#) - IMPALA-4108: codegen all builtin aggregate functions
- [IMPALA-1616](#) - Improve the Memory Limit Exceeded error report
- [IMPALA-1654](#) - General partition exprs in DDL operations.
- [IMPALA-1702](#) - Enforce single-table consistency in query analysis
- [IMPALA-1788](#) - Fold constant expressions

- [IMPALA-2013](#) - Reintroduce steps for checking HBase health in run-hbase.sh
- [IMPALA-2057](#) - Better error message for incorrect avro decimal column declaration
- [IMPALA-2521](#) - Add clustered hint to insert statements
- [IMPALA-2523](#) - Make HdfsTableSink aware of clustered input
- [IMPALA-2789](#) - More compact mem layout with null bits at the end.
- [IMPALA-2864](#) - Ensure that client connections are closed after a failed Open()
- [IMPALA-2890](#) - Support ALTER TABLE statements for Kudu tables
- [IMPALA-2905](#) - Move QueryResultSet implementations into separate module
- [IMPALA-2905](#) - Handle coordinator fragment lifecycle like all others
- [IMPALA-2916](#) - Add warning to query profile if debug build
- [IMPALA-2925](#) - Mark test_alloc_update as xfail.
- [IMPALA-3002](#) - IMPALA-1473: Cardinality observability cleanup
- [IMPALA-3125](#) - Fix assignment of equality predicates from an outer-join On-clause
- [IMPALA-3126](#) - Conservative assignment of inner-join On-clause predicates
- [IMPALA-3167](#) - Fix assignment of WHERE conjunct through grouping agg + OJ
- [IMPALA-3200](#) - move bufferpool under runtime
- [IMPALA-3201](#) - in-memory buffer pool implementation
- [IMPALA-3201](#) - reservation implementation for new buffer pool
- [IMPALA-3202](#) - refactor scratch file management into TmpFileMgr
- [IMPALA-3202](#) - DiskIoMgr improvements for new buffer pool
- [IMPALA-3211](#) - provide toolchain build id for bootstrapping
- [IMPALA-3221](#) - Copyright / license audit
- [IMPALA-3229](#) - Don't assume that AUX exists just because of shell env
- [IMPALA-3308](#) - Get expr-test passing on PPC64LE
- [IMPALA-3314](#) - Fix Avro schema loading for partitioned tables
- [IMPALA-3342](#) - Add thread counters to monitor plan fragment execution
- [IMPALA-3346](#) - DeepCopy() Kudu rows into Impala tuples
- [IMPALA-3348](#) - Avoid per-slot check vector size in KuduScanner
- [IMPALA-3398](#) - Add docs to main Impala branch
- [IMPALA-3420](#) - use gold by default
- [IMPALA-3481](#) - Use Kudu ScanToken API for scan ranges
- [IMPALA-3491](#) - Use unique db in test_scanners.py and test_aggregation.py
- [IMPALA-3491](#) - Use unique database fixture in test_partitioning.py
- [IMPALA-3491](#) - Use unique database fixture in test_insert_parquet.py
- [IMPALA-3491](#) - Use unique database fixture in test_nested_types.py
- [IMPALA-3491](#) - Use unique database fixture in test_ddl.py
- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3567](#) - Part 2, IMPALA-3899: factor out PHJ builder
- [IMPALA-3567](#) - move ExecOption profile helpers to RuntimeProfile
- [IMPALA-3586](#) - Clean up union-node.h/cc to enable improvements.
- [IMPALA-3644](#) - Make predicate order deterministic
- [IMPALA-3671](#) - Add query option to limit scratch space usage
- [IMPALA-3676](#) - Use clang as a static analysis tool
- [IMPALA-3710](#) - Kudu DML should ignore conflicts, pt2
- [IMPALA-3710](#) - Kudu DML should ignore conflicts by default
- [IMPALA-3713](#) - ,IMPALA-4439: Fix Kudu DML shell reporting
- [IMPALA-3718](#) - Add test_cancellation tests for Kudu
- [IMPALA-3718](#) - Support subset of functional-query for Kudu
- [IMPALA-3719](#) - Simplify CREATE TABLE statements with Kudu tables
- [IMPALA-3724](#) - Support Kudu non-covering range partitions

- [IMPALA-3725](#) - Support Kudu UPSERT in Impala
- [IMPALA-3726](#) - Add support for Kudu-specific column options
- [IMPALA-3739](#) - Enable stress tests on Kudu
- [IMPALA-3771](#) - Expose kudu client timeout and set default
- [IMPALA-3786](#) - Replace "cloudera" with "apache"
- [IMPALA-3786](#) - Replace "cloudera" with "apache"
- [IMPALA-3788](#) - Fix Kudu ReadMode flag checking
- [IMPALA-3788](#) - Add flag for Kudu read-your-writes
- [IMPALA-3788](#) - Support for Kudu 'read-your-writes' consistency
- [IMPALA-3808](#) - Add incubating DISCLAIMER from the Incubator Branding Guide
- [IMPALA-3809](#) - Show Kudu-specific column metadata in DESCRIBE.
- [IMPALA-3812](#) - Fix error message for unsupported types
- [IMPALA-3815](#) - clean up cross-compiled comparator
- [IMPALA-3823](#) - Add timer to measure Parquet footer reads
- [IMPALA-3838](#) - IMPALA-4495: Codegen EvalRuntimeFilters() and fixes filter stats updates
- [IMPALA-3853](#) - More RAT cleaning.
- [IMPALA-3853](#) - squeasel is MIT (and dual copyright) not Apache
- [IMPALA-3872](#) - allow providing PyPi mirror for python packages
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3884](#) - Support TYPE_TIMESTAMP for HashTableCtx::CodegenAssignNullValue()
- [IMPALA-3902](#) - Scheduler improvements for running multiple fragment instances on a single backend
- [IMPALA-3905](#) - Add single-threaded scan node.
- [IMPALA-3912](#) - test_random_rpc_timeout is flaky.
- [IMPALA-3918](#) - Fix straggler Cloudera -> ASF license headers
- [IMPALA-3920](#) - TotalStorageWaitTime counter not populated for fragments with Kudu scan node
- [IMPALA-3943](#) - Address post-merge comments.
- [IMPALA-3971](#) - IMPALA-3229: Bootstrap an Impala dev environment
- [IMPALA-3973](#) - add position and occurrence to instr()
- [IMPALA-3980](#) - qgen: re-enable Hive as a target database
- [IMPALA-3983](#) - /IMPALA-3974: Delete function jar resources after load
- [IMPALA-4000](#) - Restricted Sentry authorization for Kudu Tables
- [IMPALA-4006](#) - Fix typo in buildall.sh introduced in
- [IMPALA-4006](#) - dangerous rm -rf statements in scripts
- [IMPALA-4008](#) - Don't bake ExprContext pointers into IR code
- [IMPALA-4008](#) - don't bake in hash table and hash join pointers
- [IMPALA-4011](#) - Remove / reword messages when statestore messages are late
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4023](#) - don't attach buffered tuple streams to batches
- [IMPALA-4026](#) - Implement double-buffering for BlockingQueue
- [IMPALA-4028](#) - Improve message for improper Sentry config to make extra spaces visible
- [IMPALA-4037](#) - IMPALA-4038: fix locking during query cancellation
- [IMPALA-4042](#) - Preserve root types when substituting grouping exprs
- [IMPALA-4047](#) - Remove occurrences of 'CDH'/'cdh' from repo
- [IMPALA-4048](#) - Misc. improvements to /sessions
- [IMPALA-4054](#) - Remove serial test workarounds for IMPALA-2479.
- [IMPALA-4056](#) - Fix toSql() of DistributeParam
- [IMPALA-4058](#) - benchmark byteswap on misaligned memory
- [IMPALA-4074](#) - Configuration items duplicate in template of YARN
- [IMPALA-4080](#) - IMPALA-3638: Introduce ExecNode::Codegen()
- [IMPALA-4087](#) - TestFragmentLifecycle.test_failure_in_prepare

- [IMPALA-4091](#) - Fix backend unit to log in logs/be_tests.
- [IMPALA-4096](#) - Allow clean.sh to work from snapshots
- [IMPALA-4097](#) - Crash in kudu-scan-node-test
- [IMPALA-4098](#) - Open()/Close() partition exprs once per fragment instance.
- [IMPALA-4100](#) - 4112: Qgen: Replace EXTRACT UDF + IS [NOT] DISTINCT FROM in HiveSqlWriter
- [IMPALA-4101](#) - qgen: Hive join predicates should only contains equality functions
- [IMPALA-4102](#) - Remote Kudu reads should be reported
- [IMPALA-4104](#) - add DCHECK to ConsumeLocal() and fix tests
- [IMPALA-4110](#) - IMPALA-3853: npm.js uses Artistic License 2
- [IMPALA-4110](#) - Clean up issues found by Apache RAT
- [IMPALA-4110](#) - Apache RAT script on Impala tarballs
- [IMPALA-4111](#) - backend death tests should not produce minidumps
- [IMPALA-4116](#) - Remove 'cdh' from version string again
- [IMPALA-4116](#) - Remove 'cdh' from version string
- [IMPALA-4117](#) - Factor simple scheduler test code into own files
- [IMPALA-4118](#) - extract encryption utils from BufferedBlockMgr
- [IMPALA-4122](#) - qgen: fix bitrotted cluster unit tests
- [IMPALA-4123](#) - Fast bit unpacking
- [IMPALA-4134](#) - IMPALA-3704: Kudu INSERT improvements
- [IMPALA-4136](#) - testKudu planner test hangs if Kudu is not supported
- [IMPALA-4138](#) - Fix AcquireState() for batches that change capacity
- [IMPALA-4142](#) - qgen: Hive does not support CTEs inside sub-query blocks
- [IMPALA-4155](#) - Update default partition when table is altered
- [IMPALA-4160](#) - Remove some leftover Llama references
- [IMPALA-4160](#) - Remove Llama support.
- [IMPALA-4171](#) - Remove JAR from repo.
- [IMPALA-4180](#) - Synchronize accesses to RuntimeState::reader_contexts_
- [IMPALA-4187](#) - Switch RPC latency metrics to histograms
- [IMPALA-4188](#) - Leopard: support external Docker volumes
- [IMPALA-4193](#) - Warn when benchmarks run with sub-optimal CPU settings
- [IMPALA-4194](#) - Bump version to 2.8.0
- [IMPALA-4199](#) - Add 'SNAPSHOT' to Impala version
- [IMPALA-4204](#) - Remove KuduScanNodeTest
- [IMPALA-4205](#) - fix tmp-file-mgr-test under ASAN
- [IMPALA-4206](#) - Add column lineage regression test.
- [IMPALA-4207](#) - test infra: move Hive options from connection to cluster options
- [IMPALA-4213](#) - Fix Kudu predicates that need constant folding
- [IMPALA-4230](#) - ASF policy issues from 2.7.0 rc3.
- [IMPALA-4231](#) - fix codegen time regression
- [IMPALA-4232](#) - qgen: Hive does not support aggregates inside specific analytic clauses
- [IMPALA-4234](#) - Remove astyle config file, looks outdated.
- [IMPALA-4239](#) - fix buffer pool test failures in release build
- [IMPALA-4240](#) - qgen: Add "ParseException line missing)" to Known Errors for Hive
- [IMPALA-4241](#) - remove spurious child queries event
- [IMPALA-4253](#) - impala-server.backends.client-cache.total-clients shows negative value
- [IMPALA-4258](#) - Remove duplicated and unused test macros
- [IMPALA-4259](#) - build Impala without any test cluster setup
- [IMPALA-4260](#) - Alter table add column drops all the column stats
- [IMPALA-4266](#) - Java udf returning string can give incorrect results
- [IMPALA-4269](#) - Codegen merging exchange node

- [IMPALA-4270](#) - Gracefully fail unsupported queries with mt_dop > 0.
- [IMPALA-4274](#) - hang in buffered-block-mgr-test
- [IMPALA-4277](#) - remove unneeded LegacyTCLIService
- [IMPALA-4277](#) - remove references for unsupported s3/s3n connectors
- [IMPALA-4277](#) - allow overriding of Hive/Hadoop versions/locations
- [IMPALA-4278](#) - Don't abort Catalog startup quickly if HMS is not present
- [IMPALA-4283](#) - Ensure Kudu-specific lineage and audit behavior
- [IMPALA-4285](#) - /IMPALA-4286: Fixes for Parquet scanner with MT_DOP > 0.
- [IMPALA-4287](#) - EE tests fail to run when KUDU_IS_SUPPORTED=false
- [IMPALA-4289](#) - Mark agg slots of NDV() functions as non-nullable
- [IMPALA-4291](#) - Reduce LLVM module's preparation time
- [IMPALA-4294](#) - Make check-schema-diff.sh executable from anywhere.
- [IMPALA-4295](#) - XFAIL wildcard SSL test
- [IMPALA-4299](#) - add buildall.sh option to start test cluster
- [IMPALA-4300](#) - Speed up BloomFilter::Or with SIMD
- [IMPALA-4302](#) - ,IMPALA-2379: constant expr arg fixes
- [IMPALA-4303](#) - Do not reset() qualifier of union operands.
- [IMPALA-4309](#) - Introduce Expr rewrite phase and supporting classes
- [IMPALA-4310](#) - Make push_to_asf.py respect --apache_remote
- [IMPALA-4314](#) - Standardize on MT-related data structures
- [IMPALA-4325](#) - StmtRewrite lost parentheses of CompoundPredicate
- [IMPALA-4330](#) - Fix JSON syntax in generate_metrics.py
- [IMPALA-4335](#) - Don't send 0-row batches to clients
- [IMPALA-4338](#) - test infra data migrator: include tables' primary keys in PostgreSQL
- [IMPALA-4339](#) - ensure coredumps end up in IMPALA_HOME
- [IMPALA-4340](#) - explain how to install postgresql-9.5 or higher
- [IMPALA-4343](#) - IMPALA-4354: qgen: model INSERTs; write INSERTs from query model
- [IMPALA-4348](#) - / IMPALA-4333: Improve coordinator fragment cancellation
- [IMPALA-4350](#) - Crash with vlog level 2 in hash join node
- [IMPALA-4352](#) - test infra: store Impala/Kudu primary keys in object model
- [IMPALA-4357](#) - Fix DROP TABLE to pass analysis if the table fails to load
- [IMPALA-4362](#) - Misc. fixes for PFE counters
- [IMPALA-4363](#) - Add Parquet timestamp validation
- [IMPALA-4365](#) - Enabling end-to-end tests on a remote cluster
- [IMPALA-4369](#) - Avoid DCHECK in Parquet scanner with MT_DOP > 0
- [IMPALA-4371](#) - Incorrect DCHECK-s in hdfs-parquet-table-writer
- [IMPALA-4372](#) - 'Describe formatted' returns types in upper case
- [IMPALA-4374](#) - Use new syntax for creating TPC-DS/H tables in Kudu stress test
- [IMPALA-4377](#) - Fix Java UDF-arg buffer use-after-free in UdfExecutorTest.
- [IMPALA-4379](#) - Fix and test Kudu table type checking, follow up
- [IMPALA-4379](#) - Fix and test Kudu table type checking
- [IMPALA-4380](#) - Remove 'cloudera' from hostnames in bin/generate_minidump_collection_testdata.py
- [IMPALA-4381](#) - Incorrect AVX version of BloomFilter::Or
- [IMPALA-4383](#) - Ensure plan fragment report thread is always started
- [IMPALA-4384](#) - NPE when cols list has trailing comma
- [IMPALA-4388](#) - Fix query option reset in tests
- [IMPALA-4391](#) - fix dropped statuses in scanners
- [IMPALA-4392](#) - restore PeakMemoryUsage to DataSink profiles
- [IMPALA-4397](#) - addendum: remove stray semicolon
- [IMPALA-4397](#) - IMPALA-3259: reduce codegen time and memory

- [IMPALA-4403](#) - Implement SHOW RANGE PARTITIONS for Kudu tables
- [IMPALA-4406](#) - Add cryptography export control notice
- [IMPALA-4408](#) - Omit null bytes for Kudu scans with no nullable slots.
- [IMPALA-4409](#) - respect lock order in QueryExecState::CancelInternal()
- [IMPALA-4410](#) - Safer tear-down of RuntimeState
- [IMPALA-4411](#) - Kudu inserts violate lock ordering and could deadlock
- [IMPALA-4412](#) - Per operator timing in profile summary is incorrect when mt_dop > 0
- [IMPALA-4415](#) - Fix unassigned scan range of size 1
- [IMPALA-4421](#) - Send custom cluster & process failure test results to logs/
- [IMPALA-4427](#) - leopard: make DOCKER_IMAGE_NAME required
- [IMPALA-4432](#) - Handle internal codegen disabling properly
- [IMPALA-4433](#) - Always generate testdata using the same time zone setting
- [IMPALA-4434](#) - In Python, ".split('\n') is [''], which has length 1
- [IMPALA-4435](#) - Fix in-predicate-benchmark linking by moving templates
- [IMPALA-4436](#) - StringValue::StringCompare() should match strncmp()
- [IMPALA-4437](#) - fix crash in disk-io-mgr
- [IMPALA-4437](#) - hit DCHECK in buffered-block-mgr-test
- [IMPALA-4438](#) - Serialize test_failpoints.py to reduce memory pressure
- [IMPALA-4440](#) - lineage timestamps can go backwards across daylight savings transitions
- [IMPALA-4441](#) - Divide-by-zero in RuntimeProfile::SummaryStatsCounter::SetStats
- [IMPALA-4442](#) - Fix FE ParserTests UnsatisfiedLinkError
- [IMPALA-4444](#) - Transfer row group resources to row batch on scan failure
- [IMPALA-4446](#) - expr-test fails under ASAN
- [IMPALA-4447](#) - Rein in overly broad sed that dirties the tree
- [IMPALA-4450](#) - qgen: use string concatenation operator for postgres queries
- [IMPALA-4452](#) - Always call AggFnEvaluator::Open() before AggFnEvaluator::Init()
- [IMPALA-4454](#) - test_kudu.TestShowCreateTable flaky
- [IMPALA-4455](#) - MemPoolTest.TryAllocateAligned failure: sizeof v. alignof
- [IMPALA-4458](#) - Fix resource cleanup of cancelled mt scan nodes.
- [IMPALA-4461](#) - Make sure data gets loaded for wide hbase tables.
- [IMPALA-4465](#) - Don't hold process wide lock while serializing Runtime Profile in GetRuntimeProfileStr()
- [IMPALA-4466](#) - Improve Kudu CRUD test coverage
- [IMPALA-4470](#) - Avoid creating a NumericLiteral from NaN/infinity/-0
- [IMPALA-4476](#) - Use unique_database to stop races in test_udfs.py
- [IMPALA-4477](#) - Bump Kudu version to latest master
- [IMPALA-4477](#) - Upgrade Kudu version to latest master
- [IMPALA-4477](#) - Upgrade Kudu version to latest master
- [IMPALA-4477](#) - Upgrade Kudu version to latest master
- [IMPALA-4478](#) - Initial Kudu client mem tracking for sink
- [IMPALA-4479](#) - Use correct isSet() thrift function when evaluating constant bool exprs
- [IMPALA-4480](#) - zero_length_region_ must be as aligned as max_align_t
- [IMPALA-4488](#) - HS2 GetOperationStatus() should keep session alive
- [IMPALA-4490](#) - Only generate runtime filters for hash join nodes.
- [IMPALA-4493](#) - fix string-compare-test when using clang
- [IMPALA-4494](#) - Fix crash in SimpleScheduler
- [IMPALA-4497](#) - Fix Kudu client crash w/ SASL initialization
- [IMPALA-4498](#) - crash in to_utc_timestamp/from_utc_timestamp
- [IMPALA-4502](#) - test_partition_ddl_predicates breaks on non-HDFS filesystems
- [IMPALA-4504](#) - fix races in PlanFragmentExecutor regarding status reporting
- [IMPALA-4509](#) - Initialise Sasl-specific mutex

- [IMPALA-4510](#) - Selectively filter args for metric verification tests
- [IMPALA-4511](#) - Add missing total_time_counter() to PFE::Exec()
- [IMPALA-4512](#) - Add a script that builds Impala on stock Ubuntu 14.04
- [IMPALA-4514](#) - Fix broken exhaustive builds caused by non-nullable columns
- [IMPALA-4516](#) - Don't hold process wide lock connection_to_sessions_map_lock_ while cancelling queries
- [IMPALA-4518](#) - CopyStringVal() doesn't copy null string
- [IMPALA-4519](#) - increase timeout in TestFragmentLifecycle
- [IMPALA-4522](#) - Bound Kudu client threads to avoid stress crash
- [IMPALA-4523](#) - Correct max VARCHAR size to 65535 (2¹⁶ - 1).
- [IMPALA-4525](#) - follow-on: cleanup error handling
- [IMPALA-4525](#) - fix crash when codegen mem limit exceeded
- [IMPALA-4527](#) - Columns in Kudu tables created from Impala default to "NULL"
- [IMPALA-4529](#) - speed up parsing of identifiers
- [IMPALA-4532](#) - Fix use-after-free in ProcessBuildInputAsync()
- [IMPALA-4535](#) - Remove 'auto' from parameter list
- [IMPALA-4539](#) - fix bug when scratch batch references I/O buffers
- [IMPALA-4540](#) - Function call in DCHECK crashes scheduler
- [IMPALA-4541](#) - fix test dimensions for test_codegen_mem_limit
- [IMPALA-4542](#) - Fix use-after-free in some BE tests
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots
- [IMPALA-4553](#) - ntpd must be synchronized for kudu to start.
- [IMPALA-4554](#) - fix projection of nested collections with mt_dop > 0
- [IMPALA-4557](#) - Fix flakiness with FLAGS_stress_free_pool_alloc
- [IMPALA-4561](#) - Replace DISTRIBUTE BY with PARTITION BY in CREATE TABLE
- [IMPALA-4562](#) - Fix for crash on kerberized clusters w/o Kudu support
- [IMPALA-4564](#) - ,IMPALA-4565: mt_dop fixes for old aggs and joins
- [IMPALA-4566](#) - Kudu client glog contention can cause timeouts
- [IMPALA-4567](#) - Fix test_kudu_alter_table exhaustive failures
- [IMPALA-4570](#) - shell tarball breaks with certain setuptools versions
- [IMPALA-4571](#) - Push IN predicates to Kudu
- [IMPALA-4572](#) - Run COMPUTE STATS on Parquet tables with MT_DOP=4
- [IMPALA-4574](#) - Do not treat UUID() like a constant expr
- [IMPALA-4577](#) - Adjust maximum size of row batch queue with MT_DOP
- [IMPALA-4578](#) - Pick up bound predicates for Kudu scan nodes.
- [IMPALA-4579](#) - SHOW CREATE VIEW fails for view containing a subquery
- [IMPALA-4580](#) - Fix crash with FETCH_FIRST when #rows < result cache size
- [IMPALA-4584](#) - Make alter table operations on Kudu tables synchronous
- [IMPALA-4585](#) - Allow the \$DATABASE template in the CATCH section
- [IMPALA-4586](#) - don't constant fold in backend
- [IMPALA-4592](#) - Improve error msg for non-deterministic predicates
- [IMPALA-4594](#) - WriteSlot and CodegenWriteSlot handle escaped NULL slots differently
- [IMPALA-4595](#) - Ignore discarded functions after linking
- [IMPALA-4608](#) - Fix fragment completion times for INSERTs
- [IMPALA-4609](#) - prefix thread counters in fragment profile
- [IMPALA-4613](#) - Make sure timers are finished before sending report profile
- [IMPALA-4614](#) - Set eval cost of timestamp literals.
- [IMPALA-4619](#) - Allow NULL as default value in Kudu tables
- [IMPALA-4628](#) - Disable broken kudu test to unblock GVOs
- [IMPALA-4630](#) - remove debug webpage easter egg
- [IMPALA-4633](#) - Change broken gflag default for Kudu client mem

- [IMPALA-4636](#) - Correct Suse Linux distro string
- [IMPALA-4636](#) - Add support for SLES12 for Kudu integration
- [IMPALA-4638](#) - Run queries with MT_DOP through admission control
- [IMPALA-4642](#) - Fix TestFragmentLifecycle failures; kudu test must wait
- [IMPALA-4654](#) - KuduScanner must return when ReachedLimit()
- [IMPALA-4659](#) - fuzz test fixes
- [IMPALA-4739](#) - ExprRewriter fails on HAVING clauses
- [IMPALA-4765](#) - Avoid using several loading threads on one table
- [IMPALA-4768](#) - Improve logging of table loading
- [IMPALA-3905](#) - Add single-threaded scan node
- [IMPALA-4262](#) - LZO-scanner fails when reading large index files from S3
- [IMPALA-4277](#) - Merge "build against hadoop components in different location" into cdh5-trunk
- [IMPALA-4277](#) - build against hadoop components in different location
- [IMPALA-4322](#) - test_scanners_fuzz.py hits a DCHECK
- [IMPALA-4391](#) - fix dropped status in scanners
- [OOZIE-2194](#) - oozie job -kill doesn't work with spark action
- [OOZIE-2225](#) - Fix test failures for
- [OOZIE-2225](#) - Add wild card filter for gathering jobs
- [OOZIE-2273](#) - MiniOozie does not work outside of Oozie
- [OOZIE-2430](#) - amendAdd root logger for hive,sqoop action
- [OOZIE-2430](#) - Add root logger for hive,sqoop action
- [OOZIE-2471](#) - Show child job url tab for distcp
- [OOZIE-2503](#) - show ChildJobURLs to spark action
- [OOZIE-2517](#) - Add support for startCreatedTime and endCreatedTime filters for coord and bundles
- [OOZIE-2520](#) - SortBy filter for ordering the jobs query results
- [OOZIE-2547](#) - Add mapreduce.job.cache.files to spark action
- [OOZIE-2552](#) - Update ActiveMQ version for security and other fixes
- [OOZIE-2563](#) - Pass spark-defaults.conf to spark action
- [OOZIE-2569](#) - Adding yarn-site, core-site, hdfs-site and mapred-site into spark launcher
- [OOZIE-2606](#) - Set spark.yarn.jars to fix Spark 2.0 with Oozie
- [OOZIE-2621](#) - Use hive-exec-<version>-core instead of hive-exec in oozie-core
- [OOZIE-2658](#) - -driver-class-path can overwrite the classpath in SparkMain
- [OOZIE-2678](#) - Oozie job -kill doesn't work with tez jobs
- [OOZIE-2705](#) - Oozie Spark action ignores spark.executor.extraJavaOptions and spark.driver.extraJavaOptions
- [OOZIE-2731](#) - Set yarn.nodemanager.disk-health-checker.max-disk-utilization-per-disk-percentage to a higher value in tests
- [PARQUET-358](#) - Add support for Avro's logical types API
- [PARQUET-415](#) - Fix ByteBuffer Binary serialization
- [PIG-4052](#) - TestJobControlSleep, TestInvokerSpeed are unreliable
- [SENTRY-1260](#) - Improve error handling when ArrayIndexOutOfBoundsException in PathsUpdate.parsePath causes MetastoreCacheInitializer initialization to fail
- [SENTRY-1270](#) - Improve error handling - Database with malformed URI causes NPE in HMS plugin during DDL
- [SENTRY-1489](#) - Categorize end-to-end (e2e) tests into slow and regular tests (adapting tests for timeout)
- [SENTRY-1497](#) - Create a Sentry scale test tool to add various objects and privileges into Sentry and HMS
- [SPARK-1239](#) - Improve fetching of map output statuses
- [SPARK-8425](#) - Add blacklist mechanism for task scheduling
- [SPARK-10722](#) - RDDBlockId not found in driver-heartbeater
- [SPARK-11301](#) - [SQL] Fix case sensitivity for filter on partitioned col...
- [SPARK-11327](#) - [MESOS] spark-dispatcher doesn't pass along some spark properties
- [SPARK-11507](#) - [MLLIB] Add compact in Matrices fromBreeze

- [SPARK-11515](#) - [ML] QuantileDiscretizer should take random seed
- [SPARK-11624](#) - [SPARK-11972][SQL] fix commands that need hive to exec
- [SPARK-11823](#) - Ignores HiveThriftBinaryServerSuite's test jdbc cancel
- [SPARK-12006](#) - GaussianMixture.train crashes if an initial model is not None
- [SPARK-12316](#) - Wait a minutes to avoid cycle calling
- [SPARK-12447](#) - [YARN] Only update the states when executor is successfully launched
- [SPARK-12655](#) - [GRAPHX] GraphX does not unpersist RDDs
- [SPARK-12672](#) - Revert "[STREAMING][UI] Use the uiRoot function instead of default root path to gain the streaming batch url."
- [SPARK-12672](#) - [STREAMING][UI] Use the uiRoot function instead of default root path to gain the streaming batch url
- [SPARK-12712](#) - Fix failure in ./dev/test-dependencies when run against empty .m2 cache
- [SPARK-12874](#) - [ML] ML StringIndexer does not protect itself from column name duplication
- [SPARK-13023](#) - [PROJECT INFRA][FOLLOWUP][BRANCH-1.6] Unable to check `root` module ending up failure of Python tests
- [SPARK-13023](#) - [PROJECT INFRA][BRANCH-1.6] Fix handling of root module in modules_to_test()
- [SPARK-13112](#) - [CORE] Make sure RegisterExecutorResponse arrive before LaunchTask
- [SPARK-13207](#) - [SQL][BRANCH-1.6] Make partitioning discovery ignore _SUCCESS files
- [SPARK-13327](#) - [SPARKRR] colnames()<- allows invalid column names
- [SPARK-13410](#) - [SQL] Support unionAll for DataFrames with UDT columns
- [SPARK-13439](#) - [MESOS] Document that spark.mesos.uris is comma-separated
- [SPARK-13444](#) - [MLLIB] QuantileDiscretizer chooses bad splits on large DataFrames
- [SPARK-13444](#) - Revert "[MLLIB] QuantileDiscretizer chooses bad splits on large DataFrames"
- [SPARK-13444](#) - [MLLIB] QuantileDiscretizer chooses bad splits on large DataFrames
- [SPARK-13454](#) - [SQL] Allow users to drop a table with a name starting with an underscore
- [SPARK-13465](#) - Add a task failure listener to TaskContext
- [SPARK-13473](#) - [SQL] Don't push predicate through project with nondeterministic field
- [SPARK-13474](#) - [PROJECT INFRA] Update packaging scripts to push artifacts to home.apache.org
- [SPARK-13475](#) - [TESTS][SQL] HiveCompatibilitySuite should still run in PR builder even if a PR only changes sql/core
- [SPARK-13482](#) - [MINOR][CONFIGURATION] Make consistency of the configuration named in TransportConf
- [SPARK-13519](#) - [CORE] Driver should tell Executor to stop itself when cleaning executor's state
- [SPARK-13522](#) - [CORE] Fix the exit log place for heartbeat
- [SPARK-13522](#) - [CORE] Executor should kill itself when it's unable to heartbeat to driver more than N times
- [SPARK-13566](#) - [CORE] Avoid deadlock between BlockManager and Executor Thread
- [SPARK-13599](#) - [BUILD] remove transitive groovy dependencies from spark-hive and spark-hiveserver
- [SPARK-13601](#) - [TESTS] use 1 partition in tests to avoid race conditions
- [SPARK-13601](#) - Call failure callbacks before writer.close()
- [SPARK-13631](#) - [CORE] Thread-safe getLocationWithLargestOutputs
- [SPARK-13642](#) - [YARN][1.6-BACKPORT] Properly handle signal kill in ApplicationMaster
- [SPARK-13648](#) - Add Hive Cli to classes for isolated classloader
- [SPARK-13697](#) - [PYSPARK] Fix the missing module name of TransformFunctionSerializer.loads
- [SPARK-13705](#) - [DOCS] UpdateStateByKey Operation documentation incorrectly refers to StatefulNetworkWordCount
- [SPARK-13711](#) - [CORE] Don't call SparkUncaughtExceptionHandler in AppClient as it's in driver
- [SPARK-13755](#) - Escape quotes in SQL plan visualization node labels
- [SPARK-13760](#) - Revert "[SQL] Fix BigDecimal constructor for FloatType"
- [SPARK-13760](#) - [SQL] Fix BigDecimal constructor for FloatType
- [SPARK-13772](#) - [SQL] Fix data type mismatch for decimal
- [SPARK-13803](#) - restore the changes in SPARK-3411
- [SPARK-13806](#) - [SQL] fix rounding mode of negative float/double
- [SPARK-13845](#) - [CORE][BACKPORT-1.6] Using onBlockUpdated to replace onTaskEnd avoiding driver OOM

- [SPARK-13901](#) - [CORE] correct the logDebug information when jump to the next locality level
- [SPARK-13958](#) - Executor OOM due to unbounded growth of pointer array in...
- [SPARK-14006](#) - [SPARKR] Fix SparkR lint-r test errors in branch-1.6
- [SPARK-14058](#) - [PYTHON] Incorrect docstring in Window.order
- [SPARK-14074](#) - [SPARKR] Specify commit sha1 ID when using install_github to install intr package
- [SPARK-14107](#) - [PYSPARK][ML] Add seed as named argument to GBTs in pyspark
- [SPARK-14138](#) - [SQL] Fix generated SpecificColumnarIterator code can exceed JVM size limit for cached DataFrames
- [SPARK-14149](#) - Log exceptions in tryOrIOException
- [SPARK-14159](#) - [ML] StringIndexerModel sets output column metadata incorrectly
- [SPARK-14187](#) - [MLLIB] Fix incorrect use of binarySearch in SparseMatrix
- [SPARK-14204](#) - [SQL] Register driverClass rather than user-specified class
- [SPARK-14219](#) - [GRAPHX] Fix `pickRandomVertex` not to fall into infinit...
- [SPARK-14232](#) - [WEBUI] Fix event timeline display issue when an executor is removed with a multiple line reason
- [SPARK-14243](#) - [CORE][BACKPORT-1.6] update task metrics when removing blocks
- [SPARK-14261](#) - [SQL] Memory leak in Spark Thrift Server
- [SPARK-14298](#) - [ML][MLLIB] LDA should support disable checkpoint
- [SPARK-14322](#) - [MLLIB] Use treeAggregate instead of reduce in OnlineLDAOptimizer
- [SPARK-14357](#) - [CORE] Properly handle the root cause being a commit denied exception
- [SPARK-14368](#) - [PYSPARK] Support python.spark.worker.memory with upper-case unit
- [SPARK-14454](#) - [1.6] Better exception handling while marking tasks as failed
- [SPARK-14468](#) - Always enable OutputCommitCoordinator
- [SPARK-14495](#) - [SQL][1.6] fix resolution failure of having clause with distinct aggregate function
- [SPARK-14544](#) - [SQL] improve performance of SQL UI tab
- [SPARK-14563](#) - [ML] use a random table name instead of __THIS__ in SQLTransformer
- [SPARK-14618](#) - [ML][DOC] Updated RegressionEvaluator.metricName param doc
- [SPARK-14665](#) - [ML][PYTHON] Fixed bug with StopWordsRemover default stopwords
- [SPARK-14671](#) - [ML] Pipeline setStages should handle subclasses of PipelineStage
- [SPARK-14757](#) - [SQL] Fix nullability bug in EqualNullSafe codegen
- [SPARK-14787](#) - [SQL] Upgrade Joda-Time library from 2.9 to 2.9.3
- [SPARK-14897](#) - [CORE] Upgrade Jetty to latest version of 8
- [SPARK-14965](#) - [SQL] Indicate an exception is thrown for a missing struct field
- [SPARK-15062](#) - [SQL] Backport fix list type infer serializer issue
- [SPARK-15091](#) - [SPARKR] Fix warnings and a failure in SparkR test cases with testthat version 1.0.1
- [SPARK-15209](#) - Fix display of job descriptions with single quotes in web UI timeline
- [SPARK-15223](#) - [DOCS] fix wrongly named config reference
- [SPARK-15260](#) - Atomically resize memory pools
- [SPARK-15262](#) - Synchronize block manager / scheduler executor state
- [SPARK-15395](#) - Revert "[CORE] Use getHostString to create RpcAddress (backport for 1.6)"
- [SPARK-15395](#) - [CORE] Use getHostString to create RpcAddress
- [SPARK-15528](#) - [SQL] Fix race condition in NumberConverter
- [SPARK-15541](#) - Casting ConcurrentHashMap to ConcurrentMap
- [SPARK-15541](#) - Casting ConcurrentHashMap to ConcurrentMap
- [SPARK-15601](#) - [CORE] CircularBuffer's toString() to print only the contents written if buffer isn't full
- [SPARK-15606](#) - [CORE] Use non-blocking removeExecutor call to avoid deadlocks
- [SPARK-15613](#) - [SQL] Fix incorrect days to millis conversion due to Daylight Saving Time
- [SPARK-15613](#) - Revert "[SQL] Fix incorrect days to millis conversion due to Daylight Saving Time"
- [SPARK-15613](#) - [SQL] Fix incorrect days to millis conversion due to Daylight Saving Time
- [SPARK-15723](#) - Fixed local-timezone-brittle test where short-timezone form "EST" is ...
- [SPARK-15736](#) - [CORE][BRANCH-1.6] Gracefully handle loss of DiskStore files
- [SPARK-15761](#) - [MLLIB][PYSPARK] Load ipython when default python is Python3

- [SPARK-15827](#) - [BUILD] Publish Spark's forked sbt-pom-reader to Maven Central
- [SPARK-15891](#) - Make YARN logs less noisy
- [SPARK-15891](#) - [YARN] Clean up some logging in the YARN AM.
- [SPARK-15892](#) - [ML] Backport correctly merging AFTAggregators to branch 1.6
- [SPARK-15892](#) - Revert "[ML] Incorrectly merged AFTAggregator with zero total count"
- [SPARK-15892](#) - [ML] Incorrectly merged AFTAggregator with zero total count
- [SPARK-15915](#) - [SQL] Logical plans should use subqueries eliminated plan when override sameResult.
- [SPARK-15975](#) - Fix improper Popen retcode code handling in dev/run-tests
- [SPARK-16035](#) - [PYSPARK] Fix SparseVector parser assertion for end parenthesis
- [SPARK-16044](#) - [SQL] Backport input_file_name() for data source based on NewHadoopRDD to branch 1.6
- [SPARK-16077](#) - [PYSPARK] catch the exception from pickle.whichmodule()
- [SPARK-16078](#) - [SQL] Backport: from_utc_timestamp/to_utc_timestamp should not depends on local timezone
- [SPARK-16086](#) - [SQL] fix Python UDF without arguments
- [SPARK-16148](#) - [SCHEDULER] Allow for underscores in TaskLocation in the Executor ID
- [SPARK-16173](#) - [SQL] Can't join describe() of DataFrame in Scala 2.10
- [SPARK-16182](#) - [CORE] Utils.scala -- terminateProcess() should call Process.destroyForcibly() if and only if Process.destroy() fails
- [SPARK-16214](#) - [EXAMPLES] fix the denominator of SparkPi
- [SPARK-16230](#) - [CORE] CoarseGrainedExecutorBackend to self kill if there is an exception while creating an Executor
- [SPARK-16257](#) - [BUILD] Update spark_ec2.py to support Spark 1.6.2 and 1.6.3.
- [SPARK-16313](#) - [SQL][BRANCH-1.6] Spark should not silently drop exceptions in file listing
- [SPARK-16329](#) - [SQL][BACKPORT-1.6] Star Expansion over Table Containing No Column #14040
- [SPARK-16353](#) - [BUILD][DOC] Missing javadoc options for java unidoc
- [SPARK-16372](#) - Revert "[MLLIB] Retag RDD to tallSkinnyQR of RowMatrix"
- [SPARK-16372](#) - [MLLIB] Retag RDD to tallSkinnyQR of RowMatrix
- [SPARK-16375](#) - [WEB UI] Fixed misassigned var: numCompletedTasks was assigned to numSkippedTasks
- [SPARK-16385](#) - [CORE] Catch correct exception when calling method via reflection
- [SPARK-16409](#) - [SQL] regexp_extract with optional groups causes NPE
- [SPARK-16414](#) - [YARN] Fix bugs for "Can not get user config when calling SparkHadoopUtil.get.conf on yarn cluser mode"
- [SPARK-16440](#) - [MLLIB] Destroy broadcasted variables even on driver
- [SPARK-16440](#) - [MLLIB] Undeleted broadcast variables in Word2Vec causing OoM for long runs
- [SPARK-16488](#) - Fix codegen variable namespace collision in pmod and partitionBy
- [SPARK-16489](#) - [SQL] Guard against variable reuse mistakes in expression code generation
- [SPARK-16514](#) - [SQL] Fix various regex codegen bugs
- [SPARK-16533](#) - [CORE] resolve deadlocking in driver when executors die
- [SPARK-16656](#) - [SQL][BRANCH-1.6] Try to make CreateTableAsSelectSuite more stable
- [SPARK-16664](#) - [SQL] Fix persist call on Data frames with more than 200...
- [SPARK-16664](#) - Revert "[SQL] Fix persist call on Data frames with more than 200..."
- [SPARK-16664](#) - [SQL] Fix persist call on Data frames with more than 200...
- [SPARK-16751](#) - [HOTFIX] Also update hadoop-1 deps file to reflect derby 10.12.1.1 security fix
- [SPARK-16751](#) - Upgrade derby to 10.12.1.1
- [SPARK-16796](#) - [WEB UI] Visible passwords on Spark environment page
- [SPARK-16831](#) - Revert "[PYTHON] Fixed bug in CrossValidator.avgMetrics"
- [SPARK-16831](#) - [PYTHON] Fixed bug in CrossValidator.avgMetrics
- [SPARK-16873](#) - [CORE] Fix SpillReader NPE when spillFile has no data
- [SPARK-16925](#) - Master should call schedule() after all executor exit events, not only failures
- [SPARK-16930](#) - [YARN] Fix a couple of races in cluster app initialization.
- [SPARK-16939](#) - [SQL] Fix build error by using `Tuple1` explicitly in StringFunctionsSuite
- [SPARK-16956](#) - Make ApplicationState.MAX_NUM_RETRY configurable

- [SPARK-17003](#) - [BUILD][BRANCH-1.6] release-build.sh is missing hive-thriftserver for scala 2.11
- [SPARK-17027](#) - Revert "[ML] Avoid integer overflow in PolynomialExpansion.getPolySize"
- [SPARK-17027](#) - [ML] Avoid integer overflow in PolynomialExpansion.getPolySize
- [SPARK-17038](#) - [STREAMING] fix metrics retrieval source of 'lastReceivedBatch'
- [SPARK-17102](#) - [SQL] bypass UserDefinedGenerator for json format check
- [SPARK-17245](#) - [SQL][BRANCH-1.6] Do not rely on Hive's session state to retrieve HiveConf
- [SPARK-17316](#) - [CORE] Fix the 'ask' type parameter in 'removeExecutor'
- [SPARK-17316](#) - [CORE] Make CoarseGrainedSchedulerBackend.removeExecutor non-blocking
- [SPARK-17356](#) - [SQL][1.6] Fix out of memory issue when generating JSON for TreeNode
- [SPARK-17418](#) - Prevent kinesis-asl-assembly artifacts from being published
- [SPARK-17465](#) - [SPARK CORE] Inappropriate memory management in `org.apache.spark.storage.MemoryStore` may lead to memory leak
- [SPARK-17485](#) - Prevent failed remote reads of cached blocks from failing entire job
- [SPARK-17531](#) - [BACKPORT] Don't initialize Hive Listeners for the Execution Client
- [SPARK-17547](#) - Ensure temp shuffle data file is cleaned up after error
- [SPARK-17549](#) - [SQL] Only collect table size stat in driver for cached relation.
- [SPARK-17617](#) - [SQL] Remainder(%) expression.eval returns incorrect result on double value
- [SPARK-17618](#) - Fix invalid comparisons between UnsafeRow and other row formats
- [SPARK-17623](#) - [CORE] Clarify type of TaskEndReason with a failed task.
- [SPARK-17648](#) - [CORE] TaskScheduler really needs offers to be an IndexedSeq
- [SPARK-17649](#) - [CORE] Log how many Spark events got dropped in AsynchronousListenerBus
- [SPARK-17675](#) - [CORE] Expand Blacklist for TaskSets
- [SPARK-17721](#) - [MLLIB][BACKPORT] Fix for multiplying transposed SparseMatrix with SparseVector
- [SPARK-17850](#) - [CORE] Add a flag to ignore corrupt files
- [SPARK-17884](#) - [SQL] To resolve Null pointer exception when casting from empty string to interval type
- [SPARK-18117](#) - [CORE] Add test for TaskSetBlacklist
- [SPARK-18535](#) - Redact sensitive information from Spark logs and UI
- [SPARK-18546](#) - [CORE] Fix merging shuffle spills when using encryption
- [SPARK-18547](#) - [CORE] Propagate I/O encryption key when executors register
- [SQOOP-2983](#) - OraOop export has degraded performance with wide tables
- [SQOOP-3013](#) - Configuration "tmpjars" is not checked for empty strings before passing to MR
- [SQOOP-3028](#) - Include stack trace in the logging of exceptions in ExportTool
- [SQOOP-3034](#) - HBase import should fail fast if using anything other than as-textfile
- [SQOOP-3066](#) - Introduce an option + env variable to enable/disable SQOOP-2737 feature
- [SQOOP-3069](#) - Get OracleExportTest#testUpsertTestExport in line with SQOOP-3066
- [ZOOKEEPER-1576](#) - Zookeeper cluster - failed to connect to cluster if one of the provided IPs causes java.net.UnknownHostException
- [ZOOKEEPER-1917](#) - (Apache Zookeeper logs cleartext admin passwords) to address CVE-2014-0085
- [ZOOKEEPER-2402](#) - (Document client side properties)

Issues Fixed in CDH 5.9.x

The following topics describe issues fixed in CDH 5.9.x, from newest to oldest release. You can also review [What's New In CDH 5.9.x](#) on page 1825 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.9.3

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.9.3:

- [FLUME-2798](#) - Malformed Syslog messages can lead to OutOfMemoryException
- [FLUME-3080](#) - Close failure in HDFS Sink might cause data loss
- [FLUME-3085](#) - HDFS Sink can skip flushing some BucketWriters, might lead to data loss

- [HADOOP-11400](#) - GraphiteSink does not reconnect to Graphite after 'broken pipe'
- [HADOOP-11599](#) - Client#getTimeout should use IPC_CLIENT_PING_DEFAULT when IPC_CLIENT_PING_KEY is not configured
- [HADOOP-12672](#) - RPC timeout should not override IPC ping interval
- [HADOOP-13503](#) - Improve SaslRpcClient failure logging
- [HADOOP-13988](#) - KMSSClientProvider does not work with WebHDFS and Apache Knox w/ProxyUser
- [HADOOP-14029](#) - Fix KMSSClientProvider for non-secure proxyuser use case
- [HDFS-10715](#) - NPE when applying AvailableSpaceBlockPlacementPolicy
- [HDFS-11445](#) - FSCK shows overall health status as corrupt even one replica is corrupt
- [YARN-6360](#) - Prevent FS state dump logger from cramming other log files
- [YARN-6453](#) - fairscheduler-statedump.log gets generated regardless of service
- [HBASE-15837](#) - Memstore size accounting is wrong if postBatchMutate() throws exception
- [HBASE-16630](#) - Fragmentation in long running Bucket Cache
- [HBASE-16739](#) - Timed out exception message should include encoded region name
- [HBASE-16977](#) - VerifyReplication should log a printable representation of the row keys
- [HBASE-17501](#) - guard against NPE while reading FileTrailer and HFileBlock
- [HBASE-17673](#) - Monitored RPC Handler not shown in the WebUI
- [HBASE-17688](#) - MultiRowRangeFilter not working correctly if given same start and stop RowKey
- [HBASE-17710](#) - HBase in standalone mode creates directories with 777 permission
- [HBASE-17717](#) - Explicitly use "sasl" ACL scheme for hbase superuser
- [HBASE-17731](#) - Fractional latency reporting in MultiThreadedAction
- [HBASE-17798](#) - RpcServer.Listener.Reader can abort due to CancelledKeyException
- [HBASE-17970](#) - Set yarn.app.mapreduce.am.staging-dir when starting MiniMRCluster
- [HBASE-18096](#) - Limit HFileUtil visibility and add missing annotations
- [HIVE-9481](#) - allow column list specification in INSERT statement
- [HIVE-9567](#) - JSON SerDe not escaping special chars when writing char/varchar data
- [HIVE-11141](#) - Improve RuleRegExp when the Expression node stack gets huge
- [HIVE-11418](#) - Dropping a database in an encryption zone with CASCADE and trash enabled fails
- [HIVE-11428](#) - Performance: Struct IN() clauses are extremely slow
- [HIVE-11671](#) - Optimize RuleRegExp in DPP codepath
- [HIVE-11842](#) - Improve RuleRegExp by caching some internal data structures
- [HIVE-13390](#) - Partial backport of HIVE-13390. Backported only httpclient 4.5.2 and httpcore 4.4.4 for fixing the Apache Hive SSL vulnerability bug.
- [HIVE-14178](#) - Hive::needsToCopy should reuse FileUtils::equalsFileSystem
- [HIVE-14380](#) - Queries on tables with remote HDFS paths fail in "encryption" checks.
- [HIVE-14564](#) - Column Pruning generates out of order columns in SelectOperator which cause ArrayIndexOutOfBoundsException.
- [HIVE-14943](#) - Partial Backport of HIVE-14943 - Base Implementation (of HIVE-10924)
- [HIVE-16297](#) - Improving hive logging configuration variables
- [HIVE-16394](#) - HoS does not support queue name change in middle of session
- [HIVE-16413](#) - Create table as select does not check ownership of the location
- [HIVE-16459](#) - Forward channelInactive to RpcDispatcher
- [HIVE-16593](#) - SparkClientFactory.stop may prevent JVM from exiting
- [HIVE-16646](#) - Alias in transform ... as clause shouldn't be case sensitive
- [HIVE-16660](#) - Not able to add partition for views in hive when sentry is enabled
- [HIVE-16693](#) - beeline "source" command freezes if you have a comment in it?
- [HUE-4897](#) - [core] Import document call should provide information about the import
- [HUE-5225](#) - [core] Prevent Oozie and Job Designer duplicate example documents from being installed
- [HUE-5303](#) - [editor] Avoid XSS in the viewmodel options
- [HUE-5349](#) - [search] Query definitions can include js XSS injection

- [HUE-5659](#) - [home] Ignore history dependencies when importing document from different cluster
- [HUE-5816](#) - Changing default setting as "allowed_hosts=*"
- [HUE-5850](#) - [sentry] Prevent creating roles with empty names
- [HUE-6075](#) - [oozie] Remove email body for schema 0.1
- [HUE-6109](#) - [core] Remove the restriction on Document2 invalid chars
- [HUE-6115](#) - [core] Fix document paths for names with unicode characters
- [HUE-6131](#) - [hive] Select partition values based on the actual datatypes of the partition column
- [HUE-6133](#) - [home] Stricter check for activeEntry existence
- [HUE-6133](#) - [home] Avoid search blinking
- [HUE-6133](#) - [home] Typing on the search box crashes IE 11
- [HUE-6144](#) - [oozie] Add generic XSL template to workflow graph parser
- [HUE-6161](#) - [doc2] Log failures and continue while converting documents
- [HUE-6193](#) - [converter] Retain last_executed time when creating doc2 object
- [HUE-6197](#) - [impala] Fix XSS Vulnerability in the old editors' error messages
- [HUE-6212](#) - [oozie] Prevent XSS injection from packets
- [HUE-6212](#) - [oozie] Prevent XSS injection in coordinator cron frequency field
- [HUE-6228](#) - [core] Disable touchscreen detection on Nicescroll
- [HUE-6250](#) - [frontend] Losing # fragment of full URL on login redirect
- [HUE-6261](#) - [oozie] Avoid JS error preventing workflow action status update
- [HUE-6262](#) - [core] Converter should separate history docs from saved docs
- [HUE-6263](#) - [converter] Delete Doc2 object incase of exception
- [HUE-6264](#) - [converter] Decrease memory usage for users with very high document1 objects
- [HUE-6266](#) - [converter] Remove unnecessary call to document link
- [HUE-6295](#) - [doc2] Avoid unrelated DB calls in sync_documents after import
- [HUE-6310](#) - [doc2] [doc2] Create missing doc1 links for delete and copy operations
- [HUE-6407](#) - [pig] Play button doesn't come back after killing the running pig job
- [HUE-6446](#) - [oozie] User cant edit shared coordinator or bundle
- [HUE-6604](#) - [oozie] Fix timestamp conversion to server timezone
- [HUE-6710](#) - [notebook] Raise Django 403 builtin
- [HUE-6710](#) - [notebook] Application reachable directly by users without granted access
- [IMPALA-3794](#) - Workaround for Breakpad ID conflicts
- [IMPALA-4293](#) - query profile should include error log
- [IMPALA-4383](#) - Ensure plan fragment report thread is always started
- [IMPALA-4409](#) - respect lock order in QueryExecState::CancelInternal()
- [IMPALA-4615](#) - Fix create_table.sql command order
- [IMPALA-4787](#) - Optimize APPX_MEDIAN() memory usage
- [IMPALA-5088](#) - Fix heap buffer overflow
- [IMPALA-5193](#) - Initialize decompressor before finding first tuple
- [IMPALA-5197](#) - Erroneous corrupted Parquet file message
- [IMPALA-5252](#) - Fix crash in HiveUdfCall::GetStringVal() when mem_limit exceeded
- [IMPALA-5253](#) - Use appropriate transport for StatestoreSubscriber
- [IMPALA-5355](#) - Fix the order of Sentry roles and privileges
- [IMPALA-5469](#) - Fix exception when processing catalog update
- [OOZIE-2739](#) - Remove property expansion pattern from ShellMain's log4j properties content
- [OOZIE-2816](#) - Strip out the first command word from Sqoop action if its "sqoop"
- [OOZIE-2818](#) - Can't overwrite oozie.action.max.output.data on a per-workflow basis
- [OOZIE-2844](#) - Increase stability of Oozie actions when log4j.properties is missing or not readable
- [PARQUET-389](#) - Support predicate push down on missing columns.
- [SENTRY-1422](#) - JDO deadlocks while processing grant while a background thread processes Notificationlogs
- [SENTRY-1476](#) - SentryStore is subject to JDQL injection

- [SENTRY-1505](#) - CommitContext isn't used by anything and should be removed
- [SENTRY-1515](#) - Cleanup exception handling in SentryStore
- [SENTRY-1517](#) - SentryStore should actually use function getMSentryRole to get roles
- [SENTRY-1557](#) - getRolesForGroups(),getRoleNamesForGroups() does too many trips to the the DB
- [SENTRY-1594](#) - TransactionBlock should become generic
- [SENTRY-1609](#) - DelegateSentryStore is subject to JDQL injection
- [SENTRY-1615](#) - SentryStore should not allocate empty objects that are immediately returned
- [SENTRY-1625](#) - PrivilegeOperatePersistence can use QueryParamBuilder
- [SENTRY-1714](#) - MetastorePlugin.java should quietly return from renameAuthzObject() when both paths are null
- [SENTRY-1759](#) - UpdatableCache leaks connections
- [SOLR-8836](#) - Return 400, and a SolrException when an invalid json is provided to the update handler instead of 500.
- [SOLR-9153](#) - Update Apache commons beanutils version to 1.9.2
- [SOLR-9527](#) - Improve distribution of replicas when restoring a collection
- [SOLR-9848](#) - Lower solr.cloud.wait-for-updates-with-stale-state-pause back down from 7 seconds.
- [SOLR-10076](#) - Hiding keystore and truststore passwords from /admin/info/* outputs
- [SOLR-10430](#) - Add ls command to ZkCLI for listing sub-dirs
- [SPARK-14930](#) - Race condition in CheckpointWriter.stop()
- [SPARK-16533](#) - Spark application not handling preemption messages
- [SPARK-16845](#) - org.apache.spark.sql.catalyst.expressions.GeneratedClass\$SpecificOrdering" grows beyond 64 KB
- [SPARK-16873](#) - force spill NPE
- [SPARK-17316](#) - Don't block StandaloneSchedulerBackend.executorRemoved
- [SPARK-17485](#) - Failed remote cached block reads can lead to whole job failure
- [SPARK-19019](#) - PySpark does not work with Python 3.6.0
- [SPARK-19263](#) - DAGScheduler should avoid sending conflicting task set.
- [SPARK-19537](#) - Move pendingPartitions to ShuffleMapStage.
- [SPARK-19688](#) - Spark on Yarn Credentials File set to different application directory
- [SPARK-20922](#) - Unsafe deserialization in Spark LauncherConnection
- [SQOOP-2896](#) - Sqoop exec job fails with SQLException Access denied for user
- [SQOOP-3123](#) - Import from oracle using oraoop with map-column-java to avro fails if special characters encounter in table name or column name
- [ZOOKEEPER-2040](#) - Server to log underlying cause of SASL connection problems

Issues Fixed in CDH 5.9.2

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.9.2:

- [FLUME-2171](#) - Add Interceptor to remove headers from event
- [FLUME-2812](#) - Fix semaphore leak causing java.lang.Error: Maximum permit count exceeded in MemoryChannel
- [FLUME-2889](#) - Fixes to DateTime computations
- [FLUME-2999](#) - Kafka channel and sink should enable statically assigned partition per event via header
- [FLUME-3027](#) - Change Kafka Channel to clear offsets map after commit
- [FLUME-3031](#) - Change sequence source to reset its counter for event body on channel exception
- [FLUME-3049](#) - Make HDFS sink rotate more reliably in secure mode
- [HADOOP-7930](#) - Kerberos relogin interval in UserGroupInformation should be configurable
- [HADOOP-11619](#) - FTPFileSystem should override getDefaultPort.
- [HADOOP-12655](#) - TestHttpServer.testBindAddress bind port range is wider than expected.
- [HADOOP-13433](#) - Race in UGI.reloginFromKeytab
- [HADOOP-13590](#) - Retry until TGT expires even if the UGI renewal thread encountered exception.
- [HADOOP-13627](#) - Have an explicit KerberosAuthException for UGI to throw, text from public constants
- [HADOOP-13641](#) - Update UGI#spawnAutoRenewalThreadForUserCreds to reduce indentation

- [HADOOP-13749](#) - KMSSClientProvider combined with KeyProviderCache can result in wrong UGI being used
- [HADOOP-13805](#) - UGI.getCurrentUser() fails if user does not have a keytab associated
- [HADOOP-13838](#) - KMSTokenRenewer should close providers
- [HADOOP-13953](#) - Make FTPFileSystem's data connection mode and transfer mode configurable
- [HADOOP-14003](#) - Make additional KMS tomcat settings configurable
- [HADOOP-14195](#) - CredentialProviderFactory\$getProviders is not thread-safe
- [HDFS-10763](#) - Open files can leak permanently due to inconsistent lease update
- [HDFS-11160](#) - VolumeScanner reports write-in-progress replicas as corrupt incorrectly
- [HDFS-11229](#), [HDFS-11056](#) - failed to close meta file
- [HDFS-11275](#) - Check groupEntryIndex and throw a helpful exception on failures when removing ACL.
- [HDFS-11292](#) - log lastWrittenTxId etc info in logSyncAll
- [HDFS-11306](#) - Print remaining edit logs from buffer if edit log can't be rolled
- [HDFS-11363](#) - Need more diagnosis info when seeing Slow waitForAckedSeqno.
- [HDFS-11379](#) - DFSInputStream may infinite loop requesting block locations
- [MAPREDUCE-6172](#) - TestDbClasses timeouts are too aggressive
- [MAPREDUCE-6571](#) - JobEndNotification info logs are missing in AM container syslog
- [MAPREDUCE-6763](#) - Shuffle server listen queue is too small
- [MAPREDUCE-6817](#) - The format of job start time in JHS is different from those of submit and finish time.
- [MAPREDUCE-6839](#) - TestRecovery.testCrashed failed
- [YARN-2306](#) - Add test for leakage of reservation metrics in fair scheduler.
- [YARN-2336](#) - Fair scheduler's REST API returns a missing '[' bracket JSON for deep queue tree
- [YARN-3251](#) - Fixed a deadlock in CapacityScheduler when computing absoluteMaxAvailableCapacity in LeafQueue
- [YARN-3933](#) - FairScheduler: Multiple calls to completedContainer are not safe.
- [YARN-3957](#) - FairScheduler NPE In FairSchedulerQueueInfo causing scheduler page to return 500.
- [YARN-4363](#) - In TestFairScheduler, testcase should not create FairScheduler redundantly.
- [YARN-4544](#) - All the log messages about rolling monitoring interval are shown with WARN level
- [YARN-4555](#) - TestDefaultContainerExecutor#testContainerLaunchError fails on non-english locale environment
- [YARN-5136](#) - Error in handling event type APP_ATTEMPT_REMOVED to the scheduler
- [YARN-5752](#) - TestLocalResourcesTrackerImpl#testLocalResourceCache times out
- [YARN-5859](#) - TestResourceLocalizationService#testParallelDownloadAttemptsForPublicResource sometimes fails
- [YARN-5890](#) - FairScheduler should log information about AM-resource-usage and max-AM-share for queues
- [YARN-5920](#) - Fix deadlock in TestRMHA.testTransitionedToStandbyShouldNotHang
- [YARN-6042](#) - Dump scheduler and queue state information into FairScheduler DEBUG log.
- [YARN-6151](#) - FS preemption does not consider child queues over fairshare if the parent is under.
- [YARN-6175](#) - FairScheduler: Negative vcore for resource needed to preempt.
- [YARN-6264](#) - AM not launched when a single vcore is available on the cluster.
- [HBASE-12949](#) - Scanner can be stuck in infinite loop if the HFile is corrupted
- [HBASE-15125](#) - BackportHBaseFsck's adoptHdfsOrphan function creates region with wrong end key boundary
- [HBASE-15328](#) - sanity check the redirect used to send master info requests to the embedded regionserver.
- [HBASE-15378](#) - Scanner cannot handle heartbeat message with no results
- [HBASE-15587](#) - FSTableDescriptors.getDescriptor() logs stack trace erroneously
- [HBASE-15931](#) - Add log for long-running tasks in AsyncProcess HBASE-16289 AsyncProcess stuck messages need to print region/server
- [HBASE-15955](#) - Disable action in CatalogJanitor#setEnabled should wait for active cleanup scan to finish
- [HBASE-16032](#) - Possible memory leak in StoreScanner
- [HBASE-16062](#) - Improper error handling in WAL Reader/Writer creation
- [HBASE-16146](#) - Counters are expensive.
- [HBASE-16238](#) - It's useless to catch SESSIONEXPIRED exception and retry in RecoverableZooKeeper
- [HBASE-16266](#) - Do not throw ScannerTimeoutException when catch UnknownScannerException

- [HBASE-16304](#) - HRegion#RegionScannerImpl#handleFileNotFoundException may lead to deadlock when trying to obtain write lock on updatesLock
- [HBASE-16350](#) - Undo server abort from HBASE-14968
- [HBASE-16429](#) - FSHLog: deadlock if rollWriter called when ring buffer filled with appends
- [HBASE-16460](#) - Can't rebuild the BucketAllocator's data structures when BucketCache uses FileIOEngine
- [HBASE-16604](#) - Scanner retries on IOException can cause the scans to miss data
- [HBASE-16662](#) - Fix open POODLE vulnerabilities
- [HBASE-16721](#) - Concurrency issue in WAL unflushed seqId tracking
- [HBASE-16807](#) - RegionServer will fail to report new active Hmaster until HMaster/RegionServer failover.
- [HBASE-16841](#) - Data loss in MOB files after cloning a snapshot and deleting that snapshot
- [HBASE-16931](#) - Setting cell's seqId to zero in compaction flow might cause RS down.
- [HBASE-16960](#) - RegionServer hang when aborting
- [HBASE-17020](#) - keylen in midkey() dont computed correctly
- [HBASE-17023](#) - Region left unassigned due to AM and SSH each thinking others would do the assignment work
- [HBASE-17044](#) - Fix merge failed before creating merged region leaves meta inconsistent
- [HBASE-17058](#) - Lower epsilon used for jitter verification from HBASE-15324
- [HBASE-17069](#) - RegionServer writes invalid META entries for split daughters in some circumstances
- [HBASE-17072](#) - CPU usage starts to climb up to 90-100% when using G1GC; purge ThreadLocal usage
- [HBASE-17206](#) - FSHLog may roll a new writer successfully with unflushed entries
- [HBASE-17241](#) - Avoid compacting already compacted mob files with _del files
- [HBASE-17265](#) - Region left unassigned in master failover when region failed to open
- [HBASE-17275](#) - Assign timeout may cause region to be unassigned forever
- [HBASE-17328](#) - Properly dispose of looped replication peers
- [HBASE-17381](#) - ReplicationSourceWorkerThread can die due to unhandled exceptions
- [HBASE-17409](#) - Limit jsonp callback name to prevent xss
- [HBASE-17452](#) - Failed taking snapshot - region Manifest proto-message too large
- [HBASE-17522](#) - Handle JVM throwing runtime exceptions when we ask for details on heap usage the same as a correctly returned 'undefined'.
- [HBASE-17558](#) - ZK dumping jsp should escape HTML.
- [HBASE-17561](#) - table status page should escape values that may contain arbitrary characters.
- [HBASE-17675](#) - ReplicationEndpoint should choose new sinks if a SaslException occurs
- [HIVE-7723](#) - Explain plan for complex query with lots of partitions is slow due to in-efficient collection used to find a matching ReadEntity
- [HIVE-10329](#) - Hadoop reflectionutils has issues ADDENDUM The backport only includes the ADDENDUM part which has the ReflectionUtils class.
- [HIVE-11594](#) - Analyze Table for column names with embedded spaces
- [HIVE-11849](#) - NPE in HiveHBaseTableShapshotInputFormat in query with just count(*)
- [HIVE-12179](#) - Add option to not add spark-assembly.jar to Hive classpath
- [HIVE-12349](#) - NPE in ORG SARG for IS NULL queries on Timestamp and Date column
- [HIVE-12465](#) - Hive might produce wrong results when (outer) joins are merged
- [HIVE-12619](#) - Switching the field order within an array of structs causes the query to fail
- [HIVE-12768](#) - Thread safety: binary sortable serde decimal deserialization
- [HIVE-12780](#) - Fix the output of the history command in Beeline
- [HIVE-12789](#) - Fix output twice in the history command of Beeline
- [HIVE-12976](#) - MetaStoreDirectSql doesn't batch IN lists in all cases
- [HIVE-13129](#) - CliService leaks HMS connection
- [HIVE-13149](#) - Remove some unnecessary HMS connections from HS2
- [HIVE-13240](#) - GroupByOperator: Drop the hash aggregates when closing operator
- [HIVE-13539](#) - HiveHFileOutputFormat searching the wrong directory for HFiles
- [HIVE-13864](#) - Beeline ignores the command that follows a semicolon and comment

- [HIVE-13866](#) - flatten callstack for directSQL errors
- [HIVE-13895](#) - HoS start-up overhead in yarn-client mode
- [HIVE-13904](#) - Ignore case when retrieving ColumnInfo from RowResolver
- [HIVE-13936](#) - Add streaming support for row_number
- [HIVE-14210](#) - ExecDriver should call jobclient.close() to trigger cleanup
- [HIVE-14693](#) - Some partitions will be left out when partition number is the multiple of the option hive.msck.repair.batch.size
- [HIVE-14764](#) - Enabling "hive.metastore.metrics.enabled" throws OOM in HiveMetastore
- [HIVE-14819](#) - FunctionInfo for permanent functions shows TEMPORARY FunctionType
- [HIVE-14820](#) - RPC server for spark inside HS2 is not getting server address properly
- [HIVE-14924](#) - MSCK REPAIR table with single threaded is throwing null pointer exception
- [HIVE-15282](#) - Different modification times are used when an index is built and when its staleness is checked
- [HIVE-15291](#) - Comparison of timestamp fails if only date part is provided
- [HIVE-15338](#) - Wrong result from non-vectorized DATEDIFF with scalar parameter of type DATE/TIMESTAMP
- [HIVE-15346](#) - "values temp table" should not be an input
- [HIVE-15410](#) - WebHCat supports get/set table property with its name containing period and hyphen
- [HIVE-15485](#) - Addendum toInvestigate the DoAs failure in HoS
- [HIVE-15485](#) - Investigate the DoAs failure in HoS
- [HIVE-15517](#) - NOT (x <=> y) returns NULL if x or y is NULL
- [HIVE-15551](#) - memory leak in directsql for mysql+bonecp specific initialization
- [HIVE-15572](#) - Improve the response time for query canceling when it happens during acquiring locks
- [HIVE-15735](#) - In some cases, view objects inside a view do not have parents.
- [HIVE-15782](#) - query on parquet table returns incorrect result when hive.optimize.index.filter is set to true
- [HIVE-15872](#) - The PERCENTILE_APPROX UDAF does not work with empty set
- [HIVE-15997](#) - Resource leaks when query is cancelled
- [HIVE-16019](#) - Query fails when group by/order by on same column with uppercase name
- [HIVE-16024](#) - MSCK Repair Requires nonstrict hive.mapred.mode
- [HIVE-16047](#) - Shouldn't try to get KeyProvider unless encryption is enabled
- [HIVE-16156](#) - FileSinkOperator should delete existing output target when renaming
- [HIVE-16175](#) - Possible race condition in InstanceCache
- [HUE-3079](#) - [jb] Display jobs from JHS (YARN and Spark) and their logs when not in RM
- [HUE-4546](#) - [editor] Cannot execute any Hive query: 'NoneType' object has no attribute 'update_data'
- [HUE-4969](#) - [core] Rename ini properties for sasl buffer to be standard, fixing Support hive.server2.thrift.sasl.qop="auth-conf"
- [HUE-5268](#) - [editor] Autocompleter should not suggest columns without tables
- [HUE-5281](#) - [core] Log real user real ip when using a load balancer
- [HUE-5295](#) - [desktop] Avoid microsecond comparison for last_modified field MySQL < 5.6 doesn't support microsecond precision. <https://code.djangoproject.com/ticket/19716>. Do not change the last_modified field when migrating history queries
- [HUE-5305](#) - [home] Fix empty share document modal and improve sharing UX
- [HUE-5310](#) - [search] Use Doc2 modal in search_controller
- [HUE-5390](#) - [assist] Fix for incorrect column order
- [HUE-5408](#) - [oozie] Support old docs while saving shared workflow
- [HUE-5476](#) - [core] Fix TTL is_idle middleware check
- [HUE-5482](#) - [home] Handle multiple home/trash directories by merging them into one.
- [HUE-5533](#) - [home] Improve home page load time
- [HUE-5602](#) - [jb] Make date filter more user friendly. Add start time filter in jobs page. Make start time filter compatible with python2.6.
- [HUE-5605](#) - [metadata] Do not enable if desktop auth_password is set
- [HUE-5670](#) - [doc2] Prevent exception when doc2 object is not linked to doc1

- [HUE-5679](#) - [impala] Do not truncate the last part of rows in result downloads
- [HUE-5684](#) - [oozie] Prevent page break when hiding workflow graph. Hide workflow graph when node count > 30
- [HUE-5714](#) - [hive] Close SQL canary query "Select "Hello World""
- [HUE-5717](#) - [backend] Some operating system incorrectly detect javascript mime-type as text/x-js instead of application/javascript
- [HUE-5722](#) - [core] Avoid query redaction when string is None
- [HUE-5742](#) - [core] Set default schema to public. Allow user to provide schema name for database via ini
- [HUE-5756](#) - [doc2] Workaround for improving the query history search time
- [HUE-5758](#) - [oozie] Fix parsing nodes from XML definition
- [HUE-5769](#) - [oozie] Remove mandatory inclusion of Kill row in the workflow dashboard graph
- [HUE-5823](#) - [editor] Cancel running doc search requests when the query has changed
- [HUE-5958](#) - [pig] Fix unicode errors when handling exceptions
- [HUE-5962](#) - [hiveserver2] Update HiveServerClient user object when opening session
- [HUE-5984](#) - [search] Escaping corrupts link-meta for building external links in grid dashboard
- [HUE-6075](#) - [oozie] Remove email body while displaying external graphs in dashboard
- [HUE-6090](#) - Hue to do a keep alive on idle sessions to HiveServer2
- [HUE-6103](#) - [fb] Log filesystem initialization exceptions
- [HUE-6104](#) - [aws] Check if boto configuration section exists before adding it
- [IMPALA-1702](#) - Enforce single-table consistency in query analysis.
- [IMPALA-2864](#) - Ensure that client connections are closed after a failed Open()
- [IMPALA-3167](#) - Fix assignment of WHERE conjunct through grouping agg + OJ.
- [IMPALA-3314](#) - Fix Avro schema loading for partitioned tables.
- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3884](#) - Support TYPE_TIMESTAMP for HashTableCtx::CodegenAssignNullValue()
- [IMPALA-3983](#) - [IMPALA-3974](#): Delete function jar resources after load
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4037](#) - [IMPALA-4038](#): fix locking during query cancellation
- [IMPALA-4180](#) - Synchronize accesses to RuntimeState::reader_contexts_
- [IMPALA-4260](#) - Alter table add column drops all the column stats
- [IMPALA-4263](#) - Fix wrong omission of agg/analytic hash exchanges.
- [IMPALA-4266](#) - Java udf returning string can give incorrect results
- [IMPALA-4282](#) - Remove max length check for type strings.
- [IMPALA-4291](#) - Reduce LLVM module's preparation time
- [IMPALA-4363](#) - [IMPALA-4585](#): Add Parquet timestamp validation
- [IMPALA-4391](#) - fix dropped statuses in scanners
- [IMPALA-4433](#) - Always generate testdata using the same time zone setting
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog
- [IMPALA-4488](#) - HS2 GetOperationStatus() should keep session alive
- [IMPALA-4494](#) - [IMPALA-4540](#): Fix crash in SimpleScheduler
- [IMPALA-4516](#) - Don't hold process wide lock connection_to_sessions_map_lock_ while cancelling queries
- [IMPALA-4518](#) - CopyStringVal() doesn't copy null string
- [IMPALA-4539](#) - fix bug when scratch batch references I/O buffers
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots
- [IMPALA-4579](#) - SHOW CREATE VIEW fails for view containing a subquery
- [IMPALA-4705](#) - [IMPALA-4779](#), [IMPALA-4780](#): Fix some Expr bugs with codegen
- [IMPALA-4765](#) - Avoid using several loading threads on one table.
- [IMPALA-4767](#) - Workaround for HIVE-15653 to preserve table stats.
- [IMPALA-4916](#) - Fix maintenance of set of item sets in DisjointSet.
- [IMPALA-4929](#) - Safe concurrent access to IR function call graph

- [IMPALA-4995](#) - Fix integer overflow in TopNNode::PrepareForOutput
- [IMPALA-4997](#) - Fix overflows in Sorter::TupleIterator
- [IMPALA-5005](#) - Don't allow server to send SASL COMPLETE msg out of order
- [IMPALA-4391](#) - fix dropped status in scanners
- [OOZIE-2194](#) - oozie job -kill doesn't work with spark action
- [OOZIE-2243](#) - Kill Command does not kill the child job for java action
- [OOZIE-2519](#) - Oozie HA with SSL info is slightly incorrect
- [OOZIE-2584](#) - Eliminate Thread.sleep() calls in TestMemoryLocks
- [OOZIE-2659](#) - TestPauseTransitService is flaky
- [OOZIE-2678](#) - Oozie job -kill doesn't work with tez jobs
- [OOZIE-2731](#) - Set yarn.nodemanager.disk-health-checker.max-disk-utilization-per-disk-percentage to a higher value in tests
- [OOZIE-2742](#) - Unable to kill applications based on tag
- [OOZIE-2748](#) - NPE in LauncherMapper.printArgs()
- [OOZIE-2757](#) - Malformed xml is Spark action doc page
- [OOZIE-2777](#) - Config-default.xml longer than 64k results in java.io.UTFDataFormatException
- [PIG-5025](#) - Fix flaky test failures in TestLoad.java
- [SENTRY-1260](#) - Improve error handling - ArrayIndexOutOfBoundsException in PathsUpdate.parsePath can cause MetastoreCacheInitializer initialization to fail
- [SENTRY-1270](#) - Improve error handling - Database with malformed URI causes NPE in HMS plugin during DDL
- [SENTRY-1390](#) - Add test cases to ensure usability of URI privileges for HMS binding
- [SENTRY-1508](#) - MetastorePlugin.java does not handle properly initialization failure
- [SENTRY-1520](#) - Provide mechanism for triggering HMS full snapshot
- [SENTRY-1564](#) - Improve error detection and reporting in MetastoreCacheInitializer.java
- [SOLR-5776](#) - backportEnabled SSL tests can easily exhaust random generator entropy and block. Set the server side to SHA1PRNG as in Steve's original patch. Use less SSL in a test run. refactor SSLConfig so that SSLTestConfig can provide SSLContexts using a NullSecureRandom to prevent SSL tests from blocking on entropy starved machines Alternate (psuedo random) NullSecureRandom for Constants.SUN_OS replace NullSecureRandom w/ NotSecurePseudoRandom
- [SOLR-9284](#) - The HDFS BlockDirectoryCache should not let it's keysToRelease or names maps grow indefinitely.
- [SOLR-9330](#) - Fix AlreadyClosedException on admin/mbeans?stats=true
- [SOLR-9699](#) - [SOLR-4668](#): fix exception from core status in parallel with core reload
- [SOLR-9819](#) - Upgrade Apache commons-fileupload to 1.3.2, fixing a security vulnerability
- [SOLR-9836](#) - Add ability to recover from leader when index corruption is detected on SolrCore creation.
- [SOLR-9859](#) - backport of replication.properties cannot be updated after being written and neither replication.properties or index.properties are durable in the face of a crash. Don't log error on NoSuchFileException
- [SOLR-9901](#) - backport of, SOLR-9899: Implement move in HdfsDirectoryFactory. SOLR-9899: StandardDirectoryFactory should use optimizations for all FilterDirectory not just NRTCachingDirectory.
- [SOLR-10031](#) - Validation of filename params in ReplicationHandler
- [SOLR-10114](#), [SOLR-9941](#) - Reordered delete-by-query can delete or omit child documents
- [SOLR-10119](#) - TestReplicationHandler assertion fixes part of
- [SOLR-10121](#), [SOLR-10116](#) - BlockCache corruption with high concurrency
- [SOLR-10360](#) - Remove an extra space from Hadoop distcp cmd used by Solr backup/restore
- [SPARK-12241](#) - [YARN] Improve failure reporting in Yarn client obtainTokenForHBase()
- [SPARK-12523](#) - [YARN] Support long-running of the Spark On HBase and hive meta store.
- [SPARK-13112](#) - [CORE] Make sure RegisterExecutorResponse arrive before LaunchTask
- [SPARK-13566](#) - [CORE] Avoid deadlock between BlockManager and Executor Thread
- [SPARK-13958](#) - Executor OOM due to unbounded growth of pointer array in `inâ€¦`
- [SPARK-14204](#) - [SQL] register driverClass rather than user-specified class
- [SPARK-16044](#) - [SQL] Backport input_file_name() for data source based on NewHadoopRDD to branch 1.6
- [SPARK-16230](#) - [CORE] CoarseGrainedExecutorBackend to self kill if there is an exception while creating an Executor

- [SPARK-17245](#) - [SQL][BRANCH-1.6] Do not rely on Hive's session state to retrieve HiveConf
- [SPARK-17465](#) - [SPARK CORE] Inappropriate memory management in `org.apache.spark.storage.MemoryStore` may lead to memory leak
- [SPARK-18750](#) - [YARN] Avoid using "mapValues" when allocating containers.
- [SPARK-19178](#) - [SQL][Backport-to-1.6] convert string of large numbers to int should return null
- [SQOOP-2349](#) - Add command line option for setting transaction isolation levels for metadata queries
- [SQOOP-2896](#) - Sqoop exec job fails with SQLException Access denied for user
- [SQOOP-2909](#) - Oracle related ImportTest fails after SQOOP-2737
- [SQOOP-2911](#) - Fix failing HCatalogExportTest caused by SQOOP-2863
- [SQOOP-2950](#) - Sqoop trunk has consistent UT failures - need fixing
- [SQOOP-2990](#) - Sqoop(oracle) export [updateTableToOracle] with "--update-mode allowinsert" : app fails with java.sql.SQLException: Missing IN or OUT parameter at index
- [SQOOP-3013](#) - Configuration "tmpjars" is not checked for empty strings before passing to MR
- [SQOOP-3028](#) - Include stack trace in the logging of exceptions in ExportTool
- [SQOOP-3053](#) - Create a cmd line argument for sqoop.throwOnError and use it through SqoopOptions
- [SQOOP-3055](#) - Fixing MySQL tests failing due to ignored test inputs/configuration
- [SQOOP-3057](#) - Fixing 3rd party Oracle tests failing due to invalid case of column names
- [SQOOP-3066](#) - Introduce an option + env variable to enable/disable SQOOP-2737 feature
- [SQOOP-3068](#) - Enhance error (tool.ImportTool: Encountered IOException running import job: java.io.IOException: Expected schema) to suggest workaround
- [SQOOP-3069](#) - Get OracleExportTest#testUpsertTestExport in line with SQOOP-3066
- [SQOOP-3071](#) - Fix OracleManager to apply localTimeZone correctly in case of Date objects too
- [SQOOP-3072](#) - Reenable escaping in ImportTest#testProductWithWhiteSpaceImport for proper execution
- [SQOOP-3081](#) - use OracleEscapeUtils.escapeIdentifier in OracleUpsertOutputFormat instead of inline appending quotes
- [SQOOP-3124](#) - Fix ordering in column list query of PostgreSQL connector to reflect the logical order instead of adhoc ordering
- [SQOOP-3140](#) - Removing deprecated mapred.map.max.attempts, mapred.reduce.max.attempts entries and using the new constants directly from Hadoop instead
- [SQOOP-3159](#) - Sqoop (export + --table) with Oracle table_name having '\$' fails with error

Issues Fixed in CDH 5.9.1

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.9.1:

- [AVRO-1943](#) - Fix test TestNettyServerWithCompression.testConnectionsCount
- [CRUNCH-592](#) - Job fails for null ByteBuffer value in Avro tables
- [FLUME-2797](#) - Use SourceCounter for SyslogTcpSource
- [FLUME-2844](#) - SpillableMemoryChannel must start ChannelCounter
- [FLUME-2982](#) - Add localhost escape sequence to HDFS sink
- [FLUME-3020](#) - Improve HDFS Sink escape sequence substitution
- [HADOOP-10300](#) - Allowed deferred sending of call responses
- [HADOOP-12453](#) - Support decoding KMS Delegation Token with its own Identifier
- [HADOOP-12483](#) - Maintain wrapped SASL ordering for postponed IPC responses
- [HADOOP-12973](#) - Make DU pluggable
- [HADOOP-12974](#) - Create a CachingGetSpaceUsed implementation that uses df
- [HADOOP-12975](#) - Add jitter to CachingGetSpaceUsed's thread
- [HADOOP-13034](#) - Log message about input options in distcp lacks some items
- [HADOOP-13072](#) - WindowsGetSpaceUsed constructor should be public
- [HADOOP-13317](#) - Add logs to KMS server-side to improve supportability
- [HADOOP-13353](#) - LdapGroupsMapping getPassword should not return null when IOException throws

- [HADOOP-13526](#) - Add detailed logging in KMS for the authentication failure of proxy user
- [HADOOP-13558](#) - UserGroupInformation created from a Subject incorrectly tries to renew the Kerberos ticket
- [HADOOP-13579](#) - Fix source-level compatibility after HADOOP-11252
- [HADOOP-13638](#) - KMS should set UGI's Configuration object properly
- [HADOOP-13669](#) - KMS Server should log exceptions before throwing
- [HADOOP-13693](#) - Remove the message about HTTP OPTIONS in SPNEGO initialization message from kms audit log
- [HDFS-4210](#) - Throw helpful exception when DNS entry for JournalNode cannot be resolved
- [HDFS-6962](#) - ACLs inheritance conflict with umaskmode
- [HDFS-7413](#) - Some unit tests should use NameNodeProtocols instead of FSNameSystem
- [HDFS-7415](#) - Move FSNameSystem.resolvePath() to FSDirectory
- [HDFS-7420](#) - Delegate permission checks to FSDirectory
- [HDFS-7463](#) - Simplify FSNamesystem#getBlockLocationsUpdateTimes
- [HDFS-7478](#) - Move org.apache.hadoop.hdfs.server.namenode.NNConf to FSNamesystem
- [HDFS-7517](#) - Remove redundant non-null checks in FSNamesystem#getBlockLocations
- [HDFS-7964](#) - Add support for async edit logging
- [HDFS-8224](#) - Schedule a block for scanning if its metadata file is corrupt
- [HDFS-8269](#) - getBlockLocations() does not resolve the .reserved path and generates incorrect edit logs when updating the atime
- [HDFS-8709](#) - Clarify automatic sync in FSEditLog#logEdit
- [HDFS-8809](#) - HDFS fsck reports under construction blocks as CORRUPT
- [HDFS-9038](#) - DFS reserved space is erroneously counted towards non-DFS used
- [HDFS-9601](#) - NNThroughputBenchmark.BlockReportStats should handle NotReplicatedYetException on adding block
- [HDFS-9630](#) - DistCp minor refactoring and clean up
- [HDFS-9638](#) - Improve DistCp Help and documentation
- [HDFS-9781](#) - FsDatasetImpl#getBlockReports can occasionally throw NullPointerException
- [HDFS-9820](#) - Improve distcp to support efficient restore to an earlier snapshot
- [HDFS-10216](#) - Distcp -diff throws exception when handling relative path
- [HDFS-10270](#) - TestJMXGet:testNameNode() fails
- [HDFS-10298](#) - Document the usage of distcp -diff option
- [HDFS-10312](#) - Large block reports can fail to decode at NameNode due to 64 MB protobuf maximum length restriction
- [HDFS-10313](#) - Distcp need to enforce the order of snapshot names passed to -diff
- [HDFS-10397](#) - Distcp should ignore -delete option if -diff option is provided instead of exiting
- [HDFS-10556](#) - DistCpOptions should be validated automatically
- [HDFS-10559](#) - DiskBalancer: Use SHA1 for Plan ID
- [HDFS-10567](#) - Improve plan command help message
- [HDFS-10609](#) - Uncaught InvalidEncryptionKeyException during pipeline recovery can abort downstream applications
- [HDFS-10641](#) - TestBlockManager#testBlockReportQueueing fails intermittently
- [HDFS-10652](#) - Add a unit test for HDFS-4660
- [HDFS-10722](#) - Fix race condition in TestEditLog#testBatchedSyncWithClosedLogs
- [HDFS-10760](#) - DataXceiver#run() should not log InvalidToken exception as an error
- [HDFS-10822](#) - Log DataNodes in the write pipeline
- [HDFS-10879](#) - TestEncryptionZonesWithKMS#testReadWrite fails intermittently
- [HDFS-10962](#) - TestRequestHedgingProxyProvider is unreliable
- [HDFS-10963](#) - Reduce log level when network topology cannot find enough datanodes
- [HDFS-11012](#) - Unnecessary INFO logging on DFSClients for InvalidToken
- [HDFS-11040](#) - Add documentation for HDFS-9820 distcp improvement
- [HDFS-11056](#) - Concurrent append and read operations lead to checksum error

- [MAPREDUCE-4784](#) - TestRecovery occasionally fails
- [MAPREDUCE-6628](#) - Potential memory leak in CryptoOutputStream
- [MAPREDUCE-6641](#) - TestTaskAttempt fails in trunk
- [MAPREDUCE-6670](#) - TestJobListCache#testEviction sometimes fails on Windows with timeout
- [MAPREDUCE-6718](#) - add progress log to JHS during startup
- [MAPREDUCE-6728](#) - Give fetchers hint when ShuffleHandler rejects a shuffling connection
- [MAPREDUCE-6738](#) - TestJobListCache.testAddExisting failed intermittently in slow VM testbed
- [MAPREDUCE-6771](#) - RMContainerAllocator sends container diagnostics event after corresponding completion event
- [MAPREDUCE-6798](#) - Fix intermittent failure of TestJobHistoryParsing.testJobHistoryMethods
- [YARN-2977](#) - Fixed intermittent TestNMClient failure
- [YARN-3601](#) - Fix UT TestRMFailover.testRMWebAppRedirect
- [YARN-3654](#) - ContainerLogsPage web UI should not have meta-refresh
- [YARN-3722](#) - Merge multiple TestWebAppUtils into o.a.h.yarn.webapp.util.TestWebAppUtils
- [YARN-4004](#) - container-executor should print output of docker logs if the docker container exits with non-0 exit status
- [YARN-4017](#) - container-executor overuses PATH_MAX
- [YARN-4092](#) - Fixed UI redirection to print useful messages when both RMs are in standby mode
- [YARN-4245](#) - Generalize config file handling in container-executor
- [YARN-4255](#) - container-executor does not clean up Docker operation command files
- [YARN-4820](#) - ResourceManager web redirects in HA mode drops query parameters
- [YARN-4940](#) - yarn node -list -all fail if RM starts with decommissioned node
- [YARN-5001](#) - Aggregated Logs root directory is created with wrong group if nonexistent
- [YARN-5107](#) - TestContainerMetrics fails
- [YARN-5246](#) - NMWebAppFilter web redirects drop query parameters
- [YARN-5704](#) - Provide configuration knobs to control enabling/disabling new/work in progress features in container-executor
- [YARN-5837](#) - NPE when getting node status of a decommissioned node after an RM restart
- [YARN-5862](#) - TestDiskFailures.testLocalDirsFailures failed
- [HBASE-15324](#) - Jitter can cause desiredMaxFileSize overflow in ConstantSizeRegionSplitPolicy and trigger unexpected split
- [HBASE-15430](#) - Failed taking snapshot - Manifest proto-message too large
- [HBASE-15856](#) - Do not cache unresolved addresses for connections
- [HBASE-16172](#) - Unify the retry logic in ScannerCallableWithReplicas and RpcRetryingCallerWithReadReplicas
- [HBASE-16270](#) - Handle duplicate clearing of snapshot in region replicas
- [HBASE-16294](#) - hbck reporting "No HDFS region dir found" for replicas
- [HBASE-16345](#) - RpcRetryingCallerWithReadReplicas#call() should catch some RegionServer Exceptions
- [HBASE-16360](#) - TableMapReduceUtil addHBaseDependencyJars has the wrong class name for PrefixTreeCodec
- [HBASE-16699](#) - Overflows in AverageIntervalRateLimiter's refill() and getWaitInterval()
- [HBASE-16767](#) - Mob compaction needs to clean up files in /hbase/mobdir/.tmp and /hbase/mobdir/.tmp/.bulkload when running into IO exceptions
- [HBASE-16824](#) - Writer.flush() can be called on already closed streams in WAL roll
- [HIVE-9570](#) - Investigate test failure on union_view.q
- [HIVE-10007](#) - Support qualified table name in analyze table compute statistics for columns
- [HIVE-10384](#) - BackportRetryingMetaStoreClient does not retry wrapped TTransportExceptions
- [HIVE-10728](#) - Deprecate unix_timestamp(void) and make it deterministic
- [HIVE-10965](#) - Direct SQL for stats fails in 0-column case
- [HIVE-11901](#) - StorageBasedAuthorizationProvider requires write permission on table for SELECT statements
- [HIVE-12077](#) - MSCK Repair table should fix partitions in batches
- [HIVE-12083](#) - HIVE-10965 introduces Thrift error if partNames or colNames are empty

- [HIVE-12475](#) - Parquet schema evolution within array<struct<>> does not work
- [HIVE-12646](#) - Revert "beeline and HIVE CLI do not parse"
- [HIVE-12757](#) - Fix TestCodahaleMetrics#testFileReporting
- [HIVE-12891](#) - Hive fails when java.io.tmpdir is set to a relative location
- [HIVE-13058](#) - Add session and operation_log directory deletion messages
- [HIVE-13198](#) - Authorization issues with cascading views
- [HIVE-13237](#) - Select parquet struct field with uppercase throws NPE
- [HIVE-13381](#) - Backport introduced potential differences in the q-file output which need to be investigated further
- [HIVE-13381](#) - Timestamp and date should have precedence in type hierarchy over string group
- [HIVE-13429](#) - Tool to remove dangling scratch directory
- [HIVE-13620](#) - Merge llap branch work to master
- [HIVE-13625](#) - Hive Prepared Statement when executed with escape characters in parameter fails
- [HIVE-13645](#) - Beeline needs null-guard around hiveVars and hiveConfVars read
- [HIVE-13997](#) - Insert overwrite directory does not overwrite existing files
- [HIVE-14173](#) - NPE was thrown after enabling directsql in the middle of session
- [HIVE-14205](#) - Hive does not support union type with AVRO file format
- [HIVE-14313](#) - Test failure TestMetaStoreMetrics.testConnections
- [HIVE-14383](#) - SparkClientImpl should pass principal and keytab to spark-submit instead of calling kinit explicitly
- [HIVE-14395](#) - Add the missing data files to Avro union tests (HIVE-14205 addendum)
- [HIVE-14421](#) - FS.deleteOnExit holds references to _tmp_space.db files
- [HIVE-14426](#) - Extensive logging on info level in WebHCat
- [HIVE-14436](#) - Hive 1.2.1/Hitting "ql.Driver: FAILED: IllegalArgumentException Error"
- [HIVE-14538](#) - beeline throws exceptions with parsing Hive config when using !sh statement
- [HIVE-14697](#) - Cannot access kerberized HS2 Web UI
- [HIVE-14715](#) - Hive throws NumberFormatException with query with Null value
- [HIVE-14743](#) - ArrayIndexOutOfBoundsException - HBASE-backed views' query with JOINS
- [HIVE-14762](#) - Add logging while removing scratch space
- [HIVE-14784](#) - Operation logs are disabled automatically if the parent directory does not exist
- [HIVE-14799](#) - Query operation are not thread safe during its cancellation
- [HIVE-14805](#) - Subquery inside a view will have the object in the subquery as the direct input
- [HIVE-14810](#) - Fix failing test: TestMetaStoreMetrics.testMetaDataCounts.
- [HIVE-14817](#) - Shutdown the SessionManager timeoutChecker thread properly upon shutdown
- [HIVE-14839](#) - Improve the stability of TestSessionManagerMetrics
- [HIVE-14889](#) - Beeline leaks sensitive environment variables of HiveServer2
- [HIVE-15054](#) - Hive insertion query execution fails on Hive on Spark
- [HIVE-15061](#) - Metastore types are sometimes case sensitive
- [HIVE-15090](#) - Temporary DB failure can stop ExpiredTokenRemover thread
- [HIVE-15231](#) - Query on view with CTE and alias fails with table not found error
- [HUE-4941](#) - [editor] Content Security Policy directive blocks an image when navigating on marker map
- [HUE-5041](#) - [editor] Hue export large file to HDFS doesn't work on non-default database
- [HUE-4631](#) - [home] DB transaction failing because of atomic block on home page
- [HUE-5218](#) - [search] Validate dashboard sharing works
- [HUE-5028](#) - [security] Share Oozie workflow with modify permission;however, the user can't edit the shared WF
- [HUE-5163](#) - [security] Speed up initial page rendering
- [IMPALA-3949](#) - Log error message in FileSystemUtil.copyToLocal()
- [IMPALA-4076](#) - Fix runtime filter sort compare method
- [IMPALA-4099](#) - Fix the error message while loading UDFs with no JARs
- [IMPALA-4120](#) - Incorrect results with LEAD() analytic function
- [IMPALA-4135](#) - Thrift threaded server times out connections during high load
- [IMPALA-4153](#) - Fix count(*) on all blank("") columns - test

- [IMPALA-4170](#) - Fix identifier quoting in COMPUTE INCREMENTAL STATS
- [IMPALA-4196](#) - Cross compile bit-byte-functions
- [IMPALA-4223](#) - Handle truncated file read from HDFS cache
- [IMPALA-4237](#) - Fix materialization of 4-byte decimals in data source scan node
- [IMPALA-4246](#) - SleepForMs() utility function has undefined behavior for > 1s
- [IMPALA-4301](#) - Fix IGNORE NULLS with subquery rewriting
- [IMPALA-4336](#) - Cast exprs after unnesting union operands
- [IMPALA-4387](#) - Validate decimal type in Avro file schema
- [IMPALA-4751](#) - For unknown query IDs, /query_profile_encoded?query_id=123 starts with an empty line.
- [IMPALA-4423](#) - Correct but conservative implementation of Subquery.equals().
- [OOZIE-1814](#) - Oozie should mask any passwords in logs and REST interfaces
- [OOZIE-2582](#) - Populating external child IDs for action failures
- [OOZIE-2660](#) - Create documentation for DB Dump/Load functionality
- [PIG-3807](#) - Pig creates wrong schema after dereferencing nested tuple fields with sorts
- [PIG-3818](#) - PIG-2499 is accidentally reverted
- [SENTRY-858](#) - Add a test case for "Database prefix is not honoured when executing grant statement"
- [SENTRY-1313](#) - Database prefix is not honoured when executing grant statement
- [SENTRY-1429](#) - Backport and fix conflicts in SENTRY-1454
- [SENTRY-1464](#) - Fix Sentry e2e test failure in org.apache.sentry.tests.e2e.dbprovider.TestDbUriPermissions.testAlterPartitionLocationPrivileges
- [SQLR-9310](#) - PeerSync fails on a node restart due to IndexFingerPrint mismatch
- [SPARK-12009](#) - Avoid re-allocating YARN container while driver wants to stop all executors
- [SPARK-12339](#) - Added a NullPointerException for executor stage kill from web UI
- [SPARK-12392](#) - Optimize a location order of broadcast blocks by considering preferred local hosts
- [SPARK-12941](#) - Spark-SQL JDBC Oracle dialect fails to map string datatypes to Oracle VARCHAR datatype
- [SPARK-12966](#) - ArrayType(DecimalType) support in Postgres JDBC
- [SPARK-13242](#) - codegen fallback in case-when if there many branches
- [SPARK-13328](#) - Poor read performance for broadcast variables with dynamic resource allocation
- [SPARK-16625](#) - General data types to be mapped to Oracle
- [SPARK-16711](#) - YarnShuffleService does not re-init properly on YARN rolling upgrade
- [SPARK-17171](#) - DAG lists all partitions in the graph
- [SPARK-17433](#) - YarnShuffleService does not handle moving credentials levelDb
- [SPARK-17611](#) - Make shuffle service test really test authentication
- [SPARK-17644](#) - Do not add failedStages when abortStage for fetch failure
- [SPARK-17696](#) - Partial backport to branch-1.6.
- [SQOOP-2884](#) - Document --temporary-rootdir
- [SQOOP-2915](#) - Fixing Oracle-related unit tests
- [SQOOP-2952](#) - Fix Sqoop1 (import + --hbase-bulkload) row key not added into column family
- [SQOOP-2983](#) - OraOop export has degraded performance with wide tables
- [SQOOP-2986](#) - Add validation check for --hive-import and --incremental lastmodified
- [SQOOP-3021](#) - ClassWriter fails if a column name contains a backslash character
- [SQOOP-3034](#) - HBase import should fail fast if using anything other than as-textfile

Issues Fixed in CDH 5.9.0

CDH 5.9.0 fixes the following issues.

Apache Flume

Use SourceCounter for SyslogTcpSource

Bug: [FLUME-2797](#)

SyslogTcpSource uses a deprecated class. Use the newer SouceCounter class for SyslogTcpSource, and mark SyslogTcpSouce as deprecated.

SpillableMemoryChannel must start ChannelCounter

Bug: [FLUME-2844](#)

When using SpillableMemoryChannel, a bug causes the values of all metrics of channel component monitoring system to be zero.

Add localhost escape sequence to HDFS sink

Bug: [FLUME-2982](#)

The HDFS sink should just use localhost escape sequence instead of having to pass in a header and use the host interceptor.

Apache Hadoop

When removing a stored block, use stored BlockInfo for updating UnderReplicatedBlocks

Cloudera Bug: CDH-33233

During the removal of a stored block, BlockManager UnderReplicatedBlocks needs to be updated with the stored BlockInfo instead of the incoming Block, so that the Block with the correct generation stamp is used to calculate any pending replication that needs to be triggered.

Fix a race condition in MetricsSourceAdapter.updateJmxCache

Bug: [HADOOP-11361](#)

KMS Server should log exceptions before throwing

Cloudera Bug: CDH-45378

Bug: [HADOOP-13669](#)[HADOOP-13317](#)

When KMS throws an exception, it is not logged anywhere, and the exception message can only be seen from the client side, not the stacktrace

LdapGroupsMapping getPassword should not return null when IOException throws

Bug: [HADOOP-13353](#)

When IOException throws in getPassword(), getPassword() returns a null string. This causes setConf() to throw a java.lang.NullPointerException.

Add detailed logging in KMS for the authentication failure of a proxy user

Bug: [HADOOP-13526](#)

The log message from AuthenticationFilter.java indicates that the user was successfully authenticated. However, when the filter on DelegationTokenAuthenticationFilter is called, it hits an exception and there is no log message

UserGroupInformation created from a Subject incorrectly tries to renew the Kerberos ticket

Bug: [HADOOP-13558](#)

Cloudera Bug: CDH-45245

KMS should set UGI Configuration object properly

Bug: [HADOOP-13638](#)

The Configuration object in UGI in KMS server is not initialized properly, so it does not load core-site.xml from KMSConfiguration.KMS_CONFIG_DIR.

Throw helpful exception when DNS entry for JournalNode cannot be resolved

Bug: [HDFS-4210](#)

Cloudera Bug: CDH-43377

FSDirectory and FSNameSystem issues

Bug: [HDFS-7415](#) [HDFS-7420](#) [HDFS-7463](#) [HDFS-7478](#) [HDFS-7517](#) [HDFS-8269](#)

Cloudera Bug: CDH-44736

Schedule a block for scanning if its metadata file is corrupt

Bug: [HDFS-8224](#)

Cloudera Bug: CDH-44945

A block with corrupt metadata is not scheduled for scanning by BlockPoolSliceScanner, so it is not reported as corrupt.

DiskBalancer: Use SHA1 for plan ID

Bug: [HDFS-10559](#)

SHA1 should be used as the plan ID instead of Sha512 because much shorter and easier to handle.

Improve plan command help message

Bug: [HDFS-10567](#)

The help message for the plan command needs to be clarified and corrected.

PlanCommand#getThrsoldPercentage should not use throughput value

Bug: [HDFS-10600](#)

Cloudera Bug: CDH-41858

Namenode should use loginUser(hdfs) to generateEncryptedKey

Bug: [HDFS-10643](#)

HDFS namenode login user (hdfs) should always be used when talking to KMS to generateEncryptedKey for new file creation.

DataXceiver#run() should not log InvalidToken exception as an error

Bug: [HDFS-10760](#)

Cloudera Bug: CDH-32905

When the client has an expired token and refetches a new token, the DataNode logs an error.

Log DataNodes in the write pipeline

Bug: [HDFS-10822](#)

Cloudera Bug: CDH-43898

You cannot tell which DataNodes are involved in the write pipeline. A DEBUG trace should be added to print the list of DataNodes in the pipeline.

Reduce log level when network topology cannot find enough DataNodes

Bug: [HDFS-10963](#)

Cloudera Bug: CDH-44288

The warning-level log event should be reduced to help prevent unnecessary concern about the event.

Unnecessary INFO logging on DFSclients for InvalidToken

Bug: [HDFS-11012](#)

Cloudera Bug: CDH-32905

InvalidBlockTokenException should be changed to DEBUG to help prevent unnecessary concern about the event.

Potential memory leak in CryptoOutputStream

Bug: [MAPREDUCE-6628](#)

A potential memory leak in `CryptoOutputStream.java` allocates two direct byte buffers (`inBuffer` and `outBuffer`) that are freed when `close()` method is called.

Add progress log to JHS during startup

Bug: [MAPREDUCE-6718](#)

When the JHS starts up, it initializes the internal caches and storage through the `HistoryFileManager`. If a large number of finished jobs exist, this startup phase can last minutes without logging progress.

RMContainerAllocator sends container diagnostics event after corresponding completion event

Bug: [MAPREDUCE-6771](#)

Diagnostics information may never get into `.jhist` file, so when the job completes, the diagnostics information associated with the failed task attempts is empty.

yarn node -list -all fails if ResourceManager starts with decommissioned node

Bug: [YARN-4940](#)

Apache HBase

Superuser does not consider the keytab credentials

Bug: [HBASE-15622](#)

The superuser added by default (the process running HBase) does not take in consideration the keytab credential.

Do not cache unresolved addresses for connections

Bug: [HBASE-15856](#)

During periods where DNS is not working properly, caching can occur for connections to Master or RegionServers where the initial hostname resolution and the resolution is never reattempted. This causes clients to get `UnknownHostException` for any calls.

Overflows in AverageIntervalRateLimiter's refill() and getWaitInterval()

Bug: [HBASE-16699](#)

Mob compaction needs to clean up files in `/hbase/mobdir/.tmp` and `/hbase/mobdir/.tmp/.bulkload` when running into IO exceptions

Bug: [HBASE-16767](#)

Apache Hive

WebUI Elapsed Time may mislead users to mean query run time

Bug: [HIVE-13420](#)

Cloudera Bug: CDH-39011

When running queries from Hue, WebUI's elapsed time continuously increases.

Table not found error when trying to refresh column statistics in Metastore Manager

Bug: [HIVE-10007](#)

Cloudera Bug: CDH-42873

When you try to refresh column statistics in the Metastore Manager, it fails with the following error message:

```
Error while compiling statement: FAILED: SemanticException
[Error 10001]: Line 1:157 Table not found ''
```

StorageBasedAuthorizationProvider requires write permission on table for SELECT

Bug: [HIVE-11901](#)

Cloudera Bug: CDH-43893

StorageBasedAuthorizationProvider requires write permission on table for SELECT statements.

Subquery inside a view will have the object in the subquery as the direct input

Bug: [HIVE-14805](#)

Cloudera Bug: CDH-45077

Hive query fails from HUE for some users

Bug: [HIVE-14784](#)

Cloudera Bug: CDH-44088

Operation logging is disabled automatically for the query if for some reason the parent directory (named after the Hive session ID) created when the session is established gets deleted.

Hive throws NumberFormatException with query with Null value

Bug: [HIVE-14715](#)

Cloudera Bug: CDH-44193

Hue

[editor] Improved autocompleter

Bug: [HUE-4039](#)

[editor] Improve result table scroll performance and fix header positioning

Bug: [HUE-4910](#)

[editor] Progress status and truncating warning when direct downloading results as Excel

Bug: [HUE-4438](#)

[fb] Allow to browse S3 (other filesystems)

Bug: [HUE-2915](#)

[editor] Export query result to S3

Bug: [HUE-4367](#)

[metastore] UX create table from a file, S3 or from a directory

Bug: [HUE-4425](#)

[core] Missing some security related response headers

Bug: [HUE-4372](#)

Wildcard Certificates not supported

Cloudera Bug: CDH-39117

Hue now supports wildcard certificates and certificates using Subject Alternative Name (SAN).

Cookie without HttpOnly flag set

Cloudera Bug: CDH-41607

The HttpOnly flag is now in the response header, preventing cookies from being accessed through a client side script.

Password reveal should be disabled

Cloudera Bug: CDH-41952

The Password Reveal button is now disabled on Internet Explorer.

Web server and platform version disclosure should be prevented

Cloudera Bug: CDH-41953

Web server and platform versions were removed from the response headers.

Apache Oozie

Oozie Web Console returns 500 error when Oozie server runs on JDK 8u75 or higher

Bug: [OOZIE-2533](#)

Cloudera Bug: CDH-40362

The Oozie Web Console returns a 500 error when the Oozie server is running on JDK 8u75 and higher. The Oozie server still functions, and you can use the Oozie command line, REST API, Java API, or the Hue Oozie Dashboard to review status of those jobs.

Oozie should mask any passwords in logs and REST interfaces

Bug: [OOZIE-1814](#)

The following passwords are currently visible in the instrumentation log, REST endpoints, WebUI, and CLI and should be masked:

```
javax.net.ssl.trustStorePassword
oozie.https.keystore.pass
HADOOP_CREDSTORE_PASSWORD
OOZIE_HTTPS_KEYSTORE_PASSWORD
OOZIE_HTTPS_TRUSTSTORE_PASSWORD
```

Cloudera Search

Crunchindexertool YARN job reports success even when indexing fails

Cloudera Bug: CDH-45888

When running a Spark indexer job that fails due to shards being down, Kerberos failures, or other problems), the YARN final status is still reported as SUCCEEDED when it should be failed.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.9.0:

- [FLUME-2797](#) - Use SourceCounter for SyslogTcpSource
- [FLUME-2844](#) - SpillableMemoryChannel must start ChannelCounter
- [FLUME-2909](#) - Upgrade RAT to 0.11
- [FLUME-2920](#) - Kafka Channel Should Not Commit Offsets When Stopping
- [HADOOP-11031](#) - Design Document for Credential Provider API
- [HADOOP-11149](#) - Increase the timeout of TestZKFailoverController
- [HADOOP-11180](#) - Change log message "token.Token: Cannot find class for token kind kms-dt" to debug
- [HADOOP-11814](#) - Reformat hadoop-annotations, o.a.h. classification.tools
- [HADOOP-11872](#) - hadoop dfs command prints message about using "yarn jar" on Windows (branch-2 only)
- [HADOOP-12465](#) - Incorrect Javadoc in WritableUtils.java
- [HADOOP-12537](#) - S3A to support Amazon STS temporary credentials
- [HADOOP-12548](#) - Read s3a creds from a Credential Provider
- [HADOOP-12589](#) - Fix intermittent test failure of TestCopyPreserveFlag
- [HADOOP-12613](#) - TestFind.processArguments occasionally fails
- [HADOOP-12636](#) - Prevent ServiceLoader failure init for unused FileSystems
- [HADOOP-12723](#) - S3A: Add ability to plug in any AWSCredentialsProvider
- [HADOOP-12800](#) - Copy docker directory from 2.8 to 2.7/2.6 repos to enable pre-commit Jenkins runs
- [HADOOP-12847](#) - Hadoop daemonlog should support https and SPNEGO for Kerberized cluster
- [HADOOP-12893](#) - Verify LICENSE.txt and NOTICE.txt

- [HADOOP-12928](#) - Update netty to 3.10.5.Final to sync with ZooKeeper
- [HADOOP-12991](#) - Conflicting default ports in DelegateToFileSystem
- [HADOOP-13098](#) - Dynamic LogLevel setting page should accept case-insensitive log level string
- [HADOOP-13103](#) - Group resolution from LDAP may fail on javax.naming.ServiceUnavailableException
- [HADOOP-13154](#) - S3AFileSystem printAmazonServiceException/printAmazonClientException appear copy and paste of AWS examples
- [HADOOP-13189](#) - FairCallQueue makes callQueue larger than the configured capacity
- [HADOOP-13192](#) - org.apache.hadoop.util.LineReader cannot handle multibyte delimiters correctly
- [HADOOP-13202](#) - Avoid possible overflow in org.apache.hadoop.util.bloom.BloomFilter#getNBytes
- [HADOOP-13228](#) - Add delegation token to the connection in DelegationTokenAuthenticator
- [HADOOP-13254](#) - Create framework for configurable disk checkers
- [HADOOP-13290](#) - Appropriate use of generics in FairCallQueue
- [HADOOP-13297](#) - Add missing dependency in setting maven-remote-resource-plugin to fix builds
- [HADOOP-13298](#) - Fix the leftover L&N files in hadoop-build-tools/src/main/resources/META-INF/
- [HADOOP-13316](#) - Enforce Kerberos authentication for required ops in DelegationTokenAuthenticator
- [HADOOP-13350](#) - Additional fix to LICENSE and NOTICE
- [HADOOP-13351](#) - TestDFSClientSocketSize buffer size tests are flaky
- [HADOOP-13362](#) - DefaultMetricsSystem leaks the source name when a source unregisters
- [HADOOP-13380](#) - TestBasicDiskValidator should not write data to /tmp
- [HADOOP-13395](#) - Enhance TestKMSAudit
- [HADOOP-13443](#) - KMS should check the type of underlying keyprovider of KeyProviderExtension before falling back to default
- [HADOOP-13461](#) - NPE in KeyProvider.rollNewVersion
- [HADOOP-13476](#) - CredentialProviderFactory fails at class loading from libhdfs (JNI)
- [HADOOP-13494](#) - ReconfigurableBase can log sensitive information
- [HADOOP-13526](#) - Add detailed logging in KMS for the authentication failure of proxy user
- [HADOOP-13558](#) - UserGroupInformation created from a Subject incorrectly tries to renew the Kerberos ticket
- [HADOOP-13579](#) - Fix source-level compatibility after HADOOP-11252
- [HADOOP-13638](#) - KMS should set UGI's Configuration object properly
- [HDFS-2580](#) - NameNode#main(...) can make use of GenericOptionsParser
- [HDFS-4176](#) - EditLogTailer should call rollEdits with a timeout
- [HDFS-7415](#) - Move FSNameSystem.resolvePath() to FSDirectory
- [HDFS-7420](#) - Delegate permission checks to FSDirectory
- [HDFS-7463](#) - Simplify FSNamesystem#getBlockLocationsUpdateTimes
- [HDFS-7478](#) - Move org.apache.hadoop.hdfs.server.namenode.NNConf to FSNamesystem
- [HDFS-7517](#) - Remove redundant non-null checks in FSNamesystem#getBlockLocations
- [HDFS-7934](#) - Update RollingUpgrade rollback documentation: should use bootstrapstandby for standby NN
- [HDFS-8101](#) - DFSClient use of non-constant DFSConfigKeys pulls in WebHDFS classes at runtime
- [HDFS-8224](#) - Schedule a block for scanning if its metadata file is corrupt
- [HDFS-8269](#) - getBlockLocations() does not resolve the .reserved path and generates incorrect edit logs when updating the atime
- [HDFS-8521](#) - Add VisibleForTesting annotation to BlockPoolSlice#selectReplicaToDelete
- [HDFS-8600](#) - TestWebHdfsFileSystemContract.testGetFileBlockLocations fails in branch-2.7
- [HDFS-8608](#) - Merge HDFS-7912 to trunk and branch-2 (track BlockInfo instead of Block in UnderReplicatedBlocks and PendingReplicationBlocks)
- [HDFS-8633](#) - Fix setting of dfs.datanode.readahead.bytes in hdfs-default.xml to match DFSConfigKeys
- [HDFS-8682](#) - Should not remove decommissioned node while calculating the number of live/dead decommissioned nodes
- [HDFS-8780](#) - Fetching live/dead datanode list with arg true for removeDecommissionNode, returns list with decom node
- [HDFS-9033](#) - dfsadmin -metasave prints "NaN" for cache used%

- [HDFS-9048](#) - DistCp documentation is out-of-date
- [HDFS-9413](#) - getContentSummary() on standby should throw StandbyException
- [HDFS-9530](#) - ReservedSpace is not cleared for abandoned Blocks
- [HDFS-9533](#) - seen_txid in the shared edits directory is modified during bootstrapping
- [HDFS-9601](#) - NNThroughputBenchmark.BlockReportStats should handle NotReplicatedYetException on adding block
- [HDFS-9696](#) - Garbage snapshot records linger forever
- [HDFS-9765](#) - TestBlockScanner#testVolumeliteratorWithCaching fails intermittently
- [HDFS-9781](#) - FsDatasetImpl#getBlockReports can occasionally throw NullPointerException
- [HDFS-10225](#) - DataNode hot swap drives should disallow storage type changes
- [HDFS-10245](#) - Fix the findbugs warnings in branch-2.7
- [HDFS-10335](#) - Mover\$Processor#chooseTarget() always chooses the first matching target storage group
- [HDFS-10336](#) - TestBalancer failing intermittently because of not resetting UserGroupInformation completely
- [HDFS-10347](#) - Namenode report bad block method doesn't log the bad block or datanode
- [HDFS-10458](#) - getFileEncryptionInfo should return quickly for non-encrypted cluster
- [HDFS-10474](#) - hftp copy fails when file name with Chinese+special char in branch-2
- [HDFS-10623](#) - Remove unused import of httpClient.HttpConnection from TestWebHdfsTokens
- [HDFS-10625](#) - VolumeScanner to report why a block is found bad
- [HDFS-10641](#) - TestBlockManager#testBlockReportQueueing fails intermittently
- [HDFS-10653](#) - Optimize conversion from path string to components
- [HDFS-10688](#) - BPSERVICEActor may run into a tight loop for sending block report when hitting IOException
- [HDFS-10691](#) - FileDistribution fails in hdfs oiv command due to ArrayIndexOutOfBoundsException
- [HDFS-10693](#) - metaSave should print blocks, not LightweightHashSet
- [HDFS-10703](#) - HA NameNode Web UI should show last checkpoint time
- [HDFS-10716](#) - In Balancer, the target task should be removed when its size < 0
- [HDFS-10879](#) - TestEncryptionZonesWithKMS#testReadWrite fails intermittently
- [HDFS-10962](#) - TestRequestHedgingProxyProvider is flaky
- [HDFS-10963](#) - Reduce log level when network topology cannot find enough datanodes
- [MAPREDUCE-6242](#) - Progress report log is incredibly excessive in application master
- [MAPREDUCE-6259](#) - IllegalArgument Exception due to missing job submit time
- [MAPREDUCE-6374](#) - Distributed Cache File visibility should check permission of full path
- [MAPREDUCE-6533](#) - testDetermineCacheVisibilities of TestClientDistributedCacheManager is broken
- [MAPREDUCE-6628](#) - Potential memory leak in CryptoOutputStream
- [MAPREDUCE-6633](#) - AM should retry map attempts if the reduce task encounters compression related errors
- [MAPREDUCE-6641](#) - TestTaskAttempt fails in trunk
- [MAPREDUCE-6652](#) - Add configuration property to prevent JHS from loading jobs with a task count greater than X
- [MAPREDUCE-6718](#) - Add progress log to JHS during startup
- [MAPREDUCE-6724](#) - Single shuffle to memory must not exceed Integer#MAX_VALUE
- [MAPREDUCE-6741](#) - Add MR support to redact job conf properties
- [MAPREDUCE-6751](#) - Add debug log message when splitting is not possible due to unsplitable compression
- [MAPREDUCE-6771](#) - RMContainerAllocator sends container diagnostics event after corresponding completion event
- [YARN-3212](#) - RMNode State Transition Update with DECOMMISSIONING state
- [YARN-3226](#) - UI changes for decommissioning node
- [YARN-3445](#) - Cache runningApps in RMNode for getting running apps on given NodeId
- [YARN-4556](#) - TestFifoScheduler.testResourceOverCommit fails
- [YARN-4568](#) - Fix message when NodeManager runs into errors initializing the recovery directory
- [YARN-4702](#) - FairScheduler: Allow setting maxResources for ad hoc queues
- [YARN-4940](#) - yarn node -list -all failed if RM start with decommissioned node

- [YARN-5024](#) - TestContainerResourceUsage#testUsageAfterAMRestartWithMultipleContainers random failure
- [YARN-5343](#) - TestContinuousScheduling#testSortedNodes fails intermittently
- [YARN-5434](#) - Add -client|server argument for graceful decommission
- [YARN-5483](#) - Optimize RMAAppAttempt#pullJustFinishedContainers
- [YARN-5549](#) - AMLauncher#createAMContainerLaunchContext() should not log the command to be launched indiscriminately
- [YARN-5566](#) - Client-side NM graceful decom is not triggered when jobs finish
- [YARN-5655](#) - TestContainerManagerSecurity#testNMTokens is asserting
- [HBASE-15396](#) - Enhance mapreduce.TableSplit to add encoded region name
- [HBASE-15856](#) - Don't cache unresolved addresses for connections
- [HBASE-15889](#) - String case conversions are locale-sensitive, used without locale
- [HBASE-15891](#) - Closeable resources potentially not getting closed if exception is thrown
- [HBASE-15946](#) - Eliminate possible security concerns in Store File metrics
- [HBASE-16023](#) - Fastpath for the FIFO rpcscheduler
- [HBASE-16035](#) - Nested AutoCloseables might not all get closed
- [HBASE-16096](#) - Backport. Cleanly remove replication peers from ZooKeeper
- [HBASE-16140](#) - Bump owasp.esapi from 2.1.0 to 2.1.0.1
- [HBASE-16294](#) - hbck reporting "No HDFS region dir found" for replicas
- [HBASE-16379](#) - [replication] Minor improvement to replication/copy_tables_desc.rb
- [HBASE-16450](#) - Shell tool to dump replication queues
- [HBASE-16699](#) - Overflows in AverageIntervalRateLimiter's refill() and getWaitInterval()
- [HBASE-16767](#) - Mob compaction needs to clean up files in /hbase/mobdir/.tmp and /hbase/mobdir/.tmp/.bulkload when running into IO exceptions
- [HIVE-4570](#) - Add more information to GetOperationStatus in HiveServer2 when query is still executing
- [HIVE-6758](#) - Beeline doesn't work with -e option when started in background
- [HIVE-9013](#) - Hive set command exposes metastore db password
- [HIVE-9302](#) - Beeline add commands to register local jdbc driver names and jars
- [HIVE-9570](#) - Investigate test failure on union_view
- [HIVE-9657](#) - Use new parquet Types API builder to construct data types
- [HIVE-10190](#) - CBO: AST mode checks for TABLESAMPLE with AST.toString().contains
- [HIVE-10485](#) - Create md5 UDF
- [HIVE-10624](#) - Update the initial script to make beeline bucked cli as default and allow user choose old hive cli by env
- [HIVE-10684](#) - Fix the unit test failures for HIVE-7553 after HIVE-10674 removed the binary jar files
- [HIVE-10705](#) - Update tests for HIVE-9302 after removing binaries
- [HIVE-10722](#) - External table creation with msck in Hive can create unusable partition
- [HIVE-10755](#) - Rework on HIVE-5193 to enhance the column oriented table access
- [HIVE-10824](#) - Need to update start script changes in .cmd files
- [HIVE-10904](#) - Use beeline-log4j.properties for migrated CLI [beeline-cli Branch]
- [HIVE-10965](#) - Direct SQL for stats fails in 0-column case
- [HIVE-11028](#) - Tez: table self join and join with another table fails with IndexOutOfBoundsException
- [HIVE-11226](#) - BeeLine-Cli: support hive.cli.prompt in new CLI
- [HIVE-11236](#) - BeeLine-Cli: use the same output format as old CLI in the new CLI
- [HIVE-11280](#) - Support executing script file from hdfs in new CLI [Beeline-CLI branch]
- [HIVE-11316](#) - Use datastructure that doesnt duplicate any part of string for ASTNode::toStringTree()
- [HIVE-11336](#) - Support initial file option for new CLI [beeline-cli branch]
- [HIVE-11352](#) - Avoid the double connections with 'e' option [beeline-cli branch]
- [HIVE-11375](#) - Broken processing of queries containing NOT (x IS NOT NULL and x <> 0)
- [HIVE-11490](#) - Lazily call ASTNode::toStringTree() after tree modification
- [HIVE-11624](#) - Beeline-cli: support hive.cli.print.header in new CLI [beeline-cli branch]

- [HIVE-11637](#) - Support hive.cli.print.current.db in new CLI [beeline-cli branch]
- [HIVE-11640](#) - Shell command doesn't work for new CLI [Beeline-cli branch]
- [HIVE-11717](#) - nohup mode is not supported for new hive CLI
- [HIVE-11746](#) - Connect command should not to be allowed from user [beeline-cli branch]
- [HIVE-11796](#) - CLI option is not updated when executing the initial files [beeline-cli]
- [HIVE-11943](#) - Set old CLI as the default Client when using Hive script
- [HIVE-11944](#) - Address the review items on HIVE-11778
- [HIVE-11990](#) - Loading data inpath from a temporary table dir fails on Windows
- [HIVE-12018](#) - beeline --help doesn't return to original prompt
- [HIVE-12080](#) - Support auto type widening (int->bigint & float->double) for Parquet table
- [HIVE-12083](#) - HIVE-10965 introduces thrift error if partNames or colNames are empty
- [HIVE-12215](#) - Exchange partition does not show outputs field for post/pre execute hooks
- [HIVE-12246](#) - Orc FileDump fails with Missing CLI jar
- [HIVE-12259](#) - Command containing semicolon is broken in Beeline
- [HIVE-12345](#) - Followup for HIVE-9013 : Hidden conf vars still visible through beeline
- [HIVE-12475](#) - Parquet schema evolution within array<struct<>> does not work
- [HIVE-12590](#) - Repeated UDAFs with literals can produce incorrect result
- [HIVE-12721](#) - Add UUID built in function
- [HIVE-12785](#) - View with union type and UDF to the struct is broken
- [HIVE-12834](#) - Fix to accept the arrow keys in BeeLine CLI
- [HIVE-12983](#) - Provide a builtin function to get Hive version
- [HIVE-12987](#) - Add metrics for HS2 active users and SQL operations
- [HIVE-13058](#) - Add session and operation_log directory deletion messages
- [HIVE-13093](#) - Hive metastore does not exit on start failure
- [HIVE-13151](#) - Clean up UGI objects in FileSystem cache for transactions
- [HIVE-13198](#) - Authorization issues with cascading views
- [HIVE-13237](#) - Select parquet struct field with upper case throws NPE
- [HIVE-13381](#) - Timestamp and date should have precedence in type hierarchy than string group
- [HIVE-13420](#) - Clarify HS2 WebUI Query 'Elapsed Time'
- [HIVE-13462](#) - HiveResultSetMetaData.getPrecision() fails for NULL columns
- [HIVE-13502](#) - Beeline doesnt support session parameters in JDBC URL as documentation states
- [HIVE-13620](#) - Merge llap branch work to master
- [HIVE-13625](#) - Hive Prepared Statement when executed with escape characters in parameter fails
- [HIVE-13645](#) - Beeline needs null-guard around hiveVars and hiveConfVars read
- [HIVE-13670](#) - Improve Beeline connect/reconnect semantics
- [HIVE-13783](#) - Display a secondary prompt on beeline for multi-line statements
- [HIVE-13788](#) - hive msck listpartitions need to make use of directSQL instead of datanucleus
- [HIVE-13953](#) - Issues in HiveLockObject equals method
- [HIVE-13964](#) - Add a parameter to beeline to allow a properties file to be passed in
- [HIVE-13984](#) - Use multi-threaded approach to listing files for msck
- [HIVE-13987](#) - Clarify current error shown when HS2 is down
- [HIVE-13997](#) - Insert overwrite directory doesn't overwrite existing files
- [HIVE-14001](#) - Beeline doesn't give out an error when takes either "-e" or "-f" in command instead of both
- [HIVE-14013](#) - Describe table doesn't show unicode properly
- [HIVE-14037](#) - java.lang.ClassNotFoundException for the jar in hive.reloadable.aux.jars.path in mapreduce
- [HIVE-14074](#) - RELOAD FUNCTION should update dropped functions
- [HIVE-14075](#) - BeeLine.java.orig was accidentally committed during HIVE-14001 patch
- [HIVE-14085](#) - Allow type widening primitive conversion on hive/parquet tables
- [HIVE-14090](#) - JDOExceptions thrown by the Metastore have their full stack trace returned to clients
- [HIVE-14135](#) - Beeline output not formatted correctly for large column widths

- [HIVE-14137](#) - Hive on Spark throws FileAlreadyExistsException for jobs with multiple empty tables
- [HIVE-14149](#) - Joda Time causes an AmazonS3Exception on Hadoop3.0.0
- [HIVE-14151](#) - Use of USE_DEPRECATED_CLI environment variable does not work
- [HIVE-14153](#) - Beeline: beeline history doesn't work on Hive2
- [HIVE-14207](#) - Strip HiveConf hidden params in webui conf
- [HIVE-14215](#) - Displaying inconsistent CPU usage data with MR execution engine
- [HIVE-14226](#) - Invalid check on an ASTNode#toStringTree in CalcitePlanner
- [HIVE-14267](#) - HS2 open_operations metrics not decremented when an operation gets timed out
- [HIVE-14270](#) - Write temporary data to HDFS when doing inserts on tables located on S3
- [HIVE-14294](#) - HiveSchemaConverter for Parquet doesn't translate TINYINT and SMALLINT into proper Parquet types
- [HIVE-14296](#) - Session count is not decremented when HS2 clients do not shutdown cleanly
- [HIVE-14342](#) - Beeline output is garbled when executed from a remote shell
- [HIVE-14360](#) - Starting BeeLine after using !save, there is an error logged: "Error setting configuration: conf"
- [HIVE-14383](#) - SparkClientImpl should pass principal and keytab to spark-submit instead of calling kinit explicitly
- [HIVE-14426](#) - Extensive logging on info level in WebHCat
- [HIVE-14513](#) - Enhance custom query feature in LDAP attn to support resultset of ldap groups
- [HIVE-14588](#) - Add S3 credentials to the hidden configuration variable supported on HIVE-14207
- [HIVE-14715](#) - Hive throws NumberFormatException with query with Null value
- [HIVE-14743](#) - ArrayIndexOutOfBoundsException - HBASE-backed views' query with JOINS
- [HIVE-14784](#) - Operation logs are disabled automatically if the parent directory does not exist
- [HIVE-14799](#) - Query operation not thread safe during its cancellation
- [HIVE-14805](#) - Subquery inside a view will have the object in the subquery as the direct input
- [HIVE-14848](#) - PROPOSEDS3 creds added to a hidden list by HIVE-14588 are not working on MR jobs
- [HIVE-14889](#) - Beeline leaks sensitive environment variables of HiveServer2 when you type set
- [HUE-2645](#) - [oozie] Better Spark action UX
- [HUE-2975](#) - [aws] Add S3 UploadFileHandler and implement upload to S3
- [HUE-3065](#) - [desktop] Fix test_redact_statements
- [HUE-3065](#) - [desktop] Allow searching of raw SQL in saved queries
- [HUE-3227](#) - [connector] Java snippet
- [HUE-3278](#) - [core] Migrate Google Analytics from ga.js to analytics.js and use https if possible
- [HUE-3308](#) - [spark] Migrate Livy to external repository
- [HUE-3324](#) - [editor] Create initial jison parser
- [HUE-3325](#) - [editor] Generate the parser using the makefile
- [HUE-3326](#) - [editor] Introduce the cursor and add suggestions for possible statements
- [HUE-3327](#) - [editor] Suggest tables after select
- [HUE-3328](#) - [editor] Follow case in keyword suggestions
- [HUE-3329](#) - [editor] Include database references with table suggestions
- [HUE-3348](#) - [indexer] Update deprecated field in schema.xml
- [HUE-3394](#) - [notebook] Jar submit button is always disabled
- [HUE-3647](#) - [editor] SELECT statement that starts with a comment cannot be saved to table
- [HUE-3667](#) - [sentry] Support TListSentryPrivilegesByAuthRequest v2
- [HUE-3764](#) - [editor] When exporting results to a new table, the new table is not shown
- [HUE-3765](#) - [editor] Graphing with many columns is not readable
- [HUE-3786](#) - [editor] Search box and foreach binding in the snippet DB list dropdown
- [HUE-3788](#) - [home] No way to middle click to open document in a new tab
- [HUE-3795](#) - [editor] Scroll lock a row to enable comparison with other rows
- [HUE-3797](#) - [oozie] API to create workflow for Hive documents
- [HUE-3799](#) - [editor] Impala logs have a new line
- [HUE-3831](#) - [oozie] Show operations inside the HDFS fs graph node for external workflows

- [HUE-3838](#) - [oozie] Deprecate jar_path copy and use oozie.libpath instead for actions with jars
- [HUE-3842](#) - [core] HTTP 500 while emptying Hue 3.9 trash directory
- [HUE-3885](#) - [metastore] Add check to prevent space in table name in both create table wizards
- [HUE-3896](#) - [editor] Add search on column list result set
- [HUE-3897](#) - [metastore] Add column search on the table page
- [HUE-3913](#) - [editor] Drag and drop table names is flaky
- [HUE-3921](#) - [fb] Clicking outside of an empty upload window does not close it
- [HUE-3941](#) - [search] Use sentry base solr config for examples when Sentry is on
- [HUE-3954](#) - [search] Editing record more than once will error because of old version id
- [HUE-3956](#) - [sentry] Browse to table does not load the object in v1
- [HUE-3959](#) - [core] Double blurry login popup on /security/solr page
- [HUE-3960](#) - [editor] Implement virtual rendering of the result table
- [HUE-3981](#) - [core] Restyle the login page a little
- [HUE-3997](#) - [core] Add useradmin feature that resets axes lockouts
- [HUE-3998](#) - [oozie] Blacklisting oozie prevents hue launch
- [HUE-4002](#) - [editor] Add parser support for all generic SQL test cases from the old autocompleter
- [HUE-4003](#) - [editor] Add parser support for HDFS path completion
- [HUE-4004](#) - [editor] Make all parser tests pass from the previous autocompleter
- [HUE-4008](#) - [core] Expandable and fixed close icon for error popup
- [HUE-4014](#) - [editor] Weird scroll and menu part hidden when a few rows
- [HUE-4016](#) - [editor] Add configuration option to enable the new autocompleter
- [HUE-4017](#) - [editor] Make ace editor suggestions from the new autocompleter
- [HUE-4018](#) - [editor] Error when canceling a query that is slow being submitted
- [HUE-4020](#) - [editor] Canceling query at wrong timing breaks the editor
- [HUE-4022](#) - [doc] Promote docker image to get ramped up quicker
- [HUE-4024](#) - [editor] The new autocompleter should suggest identifiers and aliases
- [HUE-4025](#) - [editor] Add keyword completion based on the current parser capabilities
- [HUE-4027](#) - [editor] New autocompleter should support functions with descriptors and types
- [HUE-4032](#) - [editor] Identify locations of databases, tables, columns and functions
- [HUE-4037](#) - [doc] Update docker image with 3.10
- [HUE-4040](#) - [editor] US map blinks when hovering texas
- [HUE-4041](#) - [editor] Auto select axes value if there is only one possible choice
- [HUE-4043](#) - [editor] Bubble up error about queries with \u2002 fails
- [HUE-4044](#) - [editor] Copy pasted query with \u2002 fails
- [HUE-4048](#) - [editor] Clicking on .. of a file settings deletes the snippet header and the row
- [HUE-4051](#) - [hive] Have lighter Impala and Hive check configs call than list DBs
- [HUE-4054](#) - [oozie] Improve of bundle inputs
- [HUE-4055](#) - [oozie] Upgrade oozie spark action to support spark-action:0.2 schema version
- [HUE-4056](#) - [home] Click on trash icons errors
- [HUE-4058](#) - [home] Rename a directory
- [HUE-4059](#) - [editor] Lost the green bar when the query finishes on query execution
- [HUE-4062](#) - [editor] Formatting action add extra spaces after group by and skip LIMIT
- [HUE-4063](#) - [home] 'done' button on wizard redirects to old home2
- [HUE-4064](#) - [editor] Format creation and update date on the table details popover
- [HUE-4066](#) - [editor] Put a toggle all columns in the result column list
- [HUE-4067](#) - [assist] Fix issue with truncated columns of the last assist entry
- [HUE-4068](#) - [metastore] Format date and file size of table stats
- [HUE-4068](#) - [metastore] Show owners and last time instead of files and size
- [HUE-4069](#) - [metastore] Add scroll to top on sample and column pages of a table
- [HUE-4070](#) - [editor] Uncheck all freezes the browser with many columns

- [HUE-4071](#) - [metastore] Table columns should sorted on their initial table order
- [HUE-4072](#) - [editor] Focus on input search box when opening result column filtering
- [HUE-4073](#) - [editor] Scroll to column in result set does not update the horizontal scroll bar
- [HUE-4074](#) - [editor] Clicking on table in assist sometimes expand or insert the name
- [HUE-4076](#) - [editor] Column value autocomplete does not filter duplicates
- [HUE-4077](#) - [editor] Decimal not supported in bar charts
- [HUE-4079](#) - [editor] Amount of 2M is missing a zero in the Y axis of bar charts
- [HUE-4080](#) - [editor] Delete queries from editor does not send to the trash
- [HUE-4081](#) - [core] Skip idle session timeout relogin popup on running jb jobs call when idle session timeout is disabled
- [HUE-4083](#) - [editor] Bar charts gets all of the same color
- [HUE-4085](#) - [editor] Fixed result col headers when cell is wide
- [HUE-4088](#) - [oozie] Make workflow action parameter dragabble
- [HUE-4089](#) - [editor] Keyboard shortcut for a new query and save query
- [HUE-4096](#) - [oozie] Fix unit tests for HUE-4054
- [HUE-4097](#) - [core] Still print error message when Oracle connector is not configured
- [HUE-4099](#) - [editor] Allow searching of raw SQL in query history
- [HUE-4100](#) - [core] First user login checks are not displayed
- [HUE-4101](#) - [editor] Hide search input when input is empty and it loses focus
- [HUE-4102](#) - [editor] Move session settings to a context panel
- [HUE-4103](#) - [core] Upgrade Font Awesome to 4.6.3
- [HUE-4104](#) - [core] Introduce Roboto as a main font
- [HUE-4105](#) - [editor] Make the editor 'editorType' a ko variable
- [HUE-4106](#) - [editor] Move schedule edition to the context panel
- [HUE-4107](#) - [desktop] Add a 'managed' field to Document2
- [HUE-4109](#) - [oozie] First pass at refactoring Coordinator editor to be pluggable
- [HUE-4110](#) - [editor] Skeleton of running schedule actions
- [HUE-4111](#) - [editor] Toggle to show or hide the column search field on result is always on
- [HUE-4112](#) - [oozie] Automatically persist the workflow of a scheduled query
- [HUE-4115](#) - [oozie] Automatically persist the coordinator of a scheduled query
- [HUE-4116](#) - [editor] Loads the coordinator of the query
- [HUE-4118](#) - [fb] Context menu can be hidden by bottom bar
- [HUE-4120](#) - [core] Update NVD3 color palette
- [HUE-4121](#) - [oozie] Improve UX of entering args to distcp action in workflow editor
- [HUE-4122](#) - [editor] Examples should have isSaved set to true
- [HUE-4123](#) - [core] Verify id parameter values in GET requests
- [HUE-4127](#) - [oozie] SLA can not be enabled for a coordinator
- [HUE-4130](#) - [core] Triple amount of lines in the /logs page
- [HUE-4131](#) - [assist] Search on the query editors should be case insensitive
- [HUE-4133](#) - [search] Queries don't handle field names that contain spaces
- [HUE-4135](#) - [editor] Automatically update on save the variables of a scheduled query
- [HUE-4136](#) - [fb] Fix normpath and parent path logic for aws fs
- [HUE-4137](#) - [editor] Open editor, create query and save it, click on Saved queries will have one item
- [HUE-4138](#) - [editor] Last modified time of a saved query is not in the correct timezone
- [HUE-4140](#) - [editor] Do not reload saved query list when opening a saved query
- [HUE-4141](#) - [oozie] Graph breaks for external workflows when there is more than one kill node
- [HUE-4142](#) - [aws] Enable S3 browse option in filebrowser
- [HUE-4143](#) - [editor] Snippet configuration layout is broken
- [HUE-4144](#) - [editor] Add search functionality to snippet DB selection
- [HUE-4147](#) - [editor] Older queries after upgrade do not have any document1

- [HUE-4149](#) - [editor] Integrated progress report of MR jobs
- [HUE-4151](#) - [oozie] Avoid dashboard page crash when coordinator does not exist in Oozie
- [HUE-4152](#) - [oozie] Managed document property blocks access to workflows of coordinators
- [HUE-4155](#) - [home] Improve document drag to select
- [HUE-4156](#) - [oozie] Do not list the managed workflows in the coordinator editor
- [HUE-4157](#) - [aws] Fix aws config_check
- [HUE-4158](#) - [oozie] Do not lose existing values when renaming variables
- [HUE-4160](#) - [editor] Do not create multiple coordinator for a new saved query
- [HUE-4161](#) - [editor] Autocomplete tables and columns in update statements
- [HUE-4163](#) - [aws] Implement S3 mkdir operation
- [HUE-4166](#) - [metastore] Sample data dialog can't display very long table name properly
- [HUE-4167](#) - [aws] Implement S3 rmtree operation
- [HUE-4170](#) - [aws] Fix S3 object timestamp
- [HUE-4173](#) - [editor] Enable autocomplete when there are backticks in table and database names
- [HUE-4174](#) - [editor] Fix nested type autocompletion of arrays, maps and structs
- [HUE-4175](#) - [editor] Improved autocompletion after WHERE
- [HUE-4176](#) - [editor] Autocomplete field values
- [HUE-4180](#) - [editor] JS error on help icon
- [HUE-4183](#) - [fb] Offer two File Browser links when S3 is configured
- [HUE-4184](#) - [editor] Set the result legend section as fixed
- [HUE-4186](#) - [editor] Create a batch_oozie connector
- [HUE-4187](#) - [editor] Integrate batch submission in the UI
- [HUE-4190](#) - [oozie] Option to return coordinator id instead of redirect in coordinator submission popup
- [HUE-4191](#) - [oozie] Prevent 500 error when listing a workflow which coordinator is none
- [HUE-4192](#) - [editor] Document2 matching query does not exist when saving
- [HUE-4193](#) - [search] Toggling on all the fields of a grid widgets breaks
- [HUE-4194](#) - [editor] Autocomplete CREATE DATABASE
- [HUE-4195](#) - [editor] Last digit of multiquery can be hidden
- [HUE-4196](#) - [editor] Catch open notebook errors
- [HUE-4197](#) - [editor] Autocomplete for DROP TABLE, DATABASE and SCHEMA
- [HUE-4198](#) - [metastore] Sample fetch breaks the page when it fails with error 1
- [HUE-4199](#) - [editor] Make the editor 'editorMode' a ko variable
- [HUE-4200](#) - [editor] Autocompletion of SHOW statements
- [HUE-4203](#) - [editor] Column search on resultset is not empty, do a scroll to any column will hide the search input
- [HUE-4204](#) - [editor] Add autocompletion support for CAST functions
- [HUE-4205](#) - [doc2] Ignore history dependencies on import
- [HUE-4206](#) - [notebook] Ignore history and other calls in notebook mode
- [HUE-4207](#) - [doc2] Delete history command to clean up old history or orphaned history
- [HUE-4210](#) - [editor] Don't revert to Hive on browser refresh of new Impala query
- [HUE-4211](#) - [editor] Autoexpand query editor sometimes bugs the result headers
- [HUE-4212](#) - [hive] Also provide if the job is started or finished
- [HUE-4216](#) - [editor] The editor type title blinks at page load and shows all the different types
- [HUE-4217](#) - [editor] Autocomplete for DESCRIBE
- [HUE-4218](#) - [oozie] Refresh parameters of a d&dropped action
- [HUE-4219](#) - [oozie] Add quick link to selected d&dropped query
- [HUE-4220](#) - [oozie] Saving a workflow make the save button blink
- [HUE-4223](#) - [oozie] No indication that \${} can be used as parameters
- [HUE-4225](#) - [editor] Colors are flipped when sorting bar results in reverse order
- [HUE-4228](#) - [search] Used to be able to toggle the field list of grid widget
- [HUE-4229](#) - [oozie] js error when selecting another query in the hive document action

- [HUE-4230](#) - [core] Finer version number logging
- [HUE-4232](#) - [notebook] Bulk operation button is gone
- [HUE-4234](#) - [editor] Improved JOIN autocompletion
- [HUE-4236](#) - [editor] Changing fields of a query we chart should reselect default values for the axis
- [HUE-4239](#) - [liboauth] Does not handle "next" urls, always sending user back to homepage
- [HUE-4240](#) - [filebrowser] Remove unused sortby param returned from listdir
- [HUE-4241](#) - [editor] Schedule are showing up in home
- [HUE-4242](#) - [editor] Add a message when query submission takes time
- [HUE-4243](#) - [aws] Only raise aws warning if default account configs are provided
- [HUE-4244](#) - [indexer] First skeleton on scalable indexer
- [HUE-4245](#) - [editor] Autocomplete subqueries in FROM and WHERE
- [HUE-4246](#) - [libsaml] SAML Integration Doesn't Check Forwarded Protocol
- [HUE-4247](#) - [batch] Support for batch pyspark or spark
- [HUE-4248](#) - [liboauth] OAuth integration Redirect URL is Assumed to be HTTP
- [HUE-4250](#) - [editor] Add configurable timeout for autocompletion API calls
- [HUE-4251](#) - [editor] Explain on incorrect query should not error
- [HUE-4262](#) - [editor] Improve performances of the table fixed header
- [HUE-4266](#) - [indexer] Add geo, host, grok, split operations
- [HUE-4268](#) - [core] Do not remove the table fixed rows on plugin redraw
- [HUE-4269](#) - [editor] Improve sample popover
- [HUE-4270](#) - [fb] Enable S3 on the file pickers
- [HUE-4271](#) - [editor] New autocompleter should support UDAFs and numeric expressions
- [HUE-4275](#) - [core] Avoid filechooser modals to scroll to the action buttons on modal open
- [HUE-4276](#) - [editor] Autocomplete around arithmetic operations and support BETWEEN, EXISTS, RLIKE and REGEX
- [HUE-4277](#) - [fb] Add endpoint that returns enabled status for each filesystem
- [HUE-4278](#) - [useradmin] Provide basic useradmin API
- [HUE-4279](#) - [editor] Shrink back the editor if the query is not big enough and it hasn't been resized manually
- [HUE-4281](#) - [editor] Fail gracefully and return an error when trying to import a wrong/corrupted JSON file
- [HUE-4285](#) - [editor] Fix Scatter plot display
- [HUE-4290](#) - [editor] Style the coordinator history on the context panel
- [HUE-4291](#) - [metastore] Do not use smallints in the create table from file wizard
- [HUE-4296](#) - [core] Need info message if kt_renewer exits due to no hue_keytab
- [HUE-4298](#) - [security] Offer ability to resize the Impersonate the user drop down in Hue security app
- [HUE-4301](#) - [editor] Autocomplete for Impala extract function
- [HUE-4302](#) - [editor] New autocompleter should suggest DISTINCT or ALL where appropriate in aggregate functions
- [HUE-4303](#) - [useradmin] Reinroduce '?' icon displaying description of each field
- [HUE-4304](#) - [core] Pie slices shrink after wrong timing of resize of the pie
- [HUE-4306](#) - [editor] The new autocompleter should complete CASE functions
- [HUE-4307](#) - [editor] Goto line number shorter conflicts with highlighting the address bar contents in Safari
- [HUE-4308](#) - [editor] UX lock rows keep adding rows when clicking on the same
- [HUE-4309](#) - [editor] UX lock rows should tooltip about clicking to lock and delete icon on hover
- [HUE-4310](#) - [fb] File chooser picker modal should have action buttons fixed to the bottom
- [HUE-4311](#) - [fb] Add boto http_socket_timeout configuration
- [HUE-4312](#) - [editor] Load query history, loses editor type when clicking on new query
- [HUE-4313](#) - [editor] The new autocompleter should accept non-reserved keywords for DBs, tables, columns etc
- [HUE-4315](#) - [editor] Error when sharing a query with R permission
- [HUE-4317](#) - [jb] Logs on attempt page do not make links to job id
- [HUE-4320](#) - [editor] New autocompleter should make suggestions based on type in value expressions
- [HUE-4322](#) - [impala] Better error message when user already downloaded all the impala resultset cache
- [HUE-4324](#) - [editor] Batch API get_jobs() should return the job ids

- [HUE-4325](#) - [editor] Toggle the execute action depending on if last execution was batch or not
- [HUE-4326](#) - [editor] Keep snippet DB list in sync like the assist
- [HUE-4327](#) - [editor] Disable batch submit by default when oozie is not setup
- [HUE-4328](#) - [editor] File chooser in 'In HDFS (large file)' always open on /
- [HUE-4329](#) - [editor] Multi query does not go past statement 1
- [HUE-4331](#) - [editor] Log call should return only the additional ones after the first call
- [HUE-4332](#) - [indexer] Add Hue logs as a supported file type
- [HUE-4334](#) - [editor] Update status to error when canceling batch job
- [HUE-4335](#) - [editor] New autocompleter should handle date types
- [HUE-4337](#) - [editor] The autocompleter should suggest backticked values when applicable
- [HUE-4340](#) - [indexer] Fix double matching and set required field default to false
- [HUE-4341](#) - [indexer] Use info level logging for morphlines, add error handling for invalid dates
- [HUE-4342](#) - [editor] DB prefix autocomplete can be lost
- [HUE-4343](#) - [editor] Improved autocompletion in select list
- [HUE-4344](#) - [fb] Buffer reading parquet file and limit in size
- [HUE-4345](#) - [aws] Fix Actions when viewing an S3 file in filebrowser
- [HUE-4347](#) - [core] Rename tests_convert.py to follow test conventions
- [HUE-4348](#) - [indexer] Refactor format types such that all information is in one place
- [HUE-4351](#) - [editor] Show number of columns in the query result
- [HUE-4352](#) - [indexer] Fix typo for grok dictionary file exists check
- [HUE-4354](#) - [core] Warn than SQLite is the cause of the 'Database locked' error
- [HUE-4355](#) - [indexer] Add i18n for arguments and extension based file type guessing
- [HUE-4356](#) - [fb] File chooser picker does not select S3 tab when opening an existing path
- [HUE-4357](#) - [oozie] Return the launcher task logs when submitting Oozie batch job
- [HUE-4358](#) - [meta] Split metadata services into two
- [HUE-4360](#) - [meta] Read configuration from a properties file
- [HUE-4363](#) - [core] Do not rely on import hadoop config before loading core
- [HUE-4366](#) - [metastore] Add an option to create a table from a file in an external location
- [HUE-4367](#) - [editor] Export query result to S3
- [HUE-4371](#) - [editor] Enable the new autocompleter by default
- [HUE-4372](#) - [core] Backport Django Security middleware or add HTTP security headers
- [HUE-4373](#) - [home] Exporting / Importing queries should handle associated workflows and coordinator
- [HUE-4374](#) - [editor] No autocomplete appearing in ORDER BY with an operation
- [HUE-4376](#) - [indexer] Added assist panel and filechooser
- [HUE-4377](#) - [editor] Autocompleter should support variable references like '\${var}' in statements
- [HUE-4378](#) - [core] File chooser with multiple options overflow the modals
- [HUE-4379](#) - [editor] Autocompleter should support CREATE TABLE statements
- [HUE-4380](#) - [assist] Only enable navigator search when configured
- [HUE-4384](#) - [indexer] Add support for combined apache log files
- [HUE-4385](#) - [indexer] Stop the indexer from dropping text fields that are under 100 characters long
- [HUE-4386](#) - [editor] Make it possible to turn autocomplete on or off
- [HUE-4390](#) - [core] Do not re-create a new user for each request in demo mode
- [HUE-4391](#) - [editor] Blacklisting Oozie should not break the editor
- [HUE-4392](#) - [fb] S3 buckets should have a different icon
- [HUE-4393](#) - [fb] S3 bucket should not have any file upload or file creation action
- [HUE-4394](#) - [fb] Home button redirects to HDFS even if on S3
- [HUE-4395](#) - [fb] S3 delete bucket implementation
- [HUE-4396](#) - [fb] Delete button should be 'delete forever' all the time
- [HUE-4397](#) - [fb] Hide upload archive in S3
- [HUE-4398](#) - [fb] API should check if bucket already exists in the region and bubble up the error

- [HUE-4399](#) - [fb] Renaming a key (non-bucket directory) creates an extraneous empty file
- [HUE-4400](#) - [fb] Disable chmod and summary for S3 for now
- [HUE-4401](#) - [fb] Do not return any date for bucket and directory objects
- [HUE-4402](#) - [fb] S3 copy action is not using the picker and fails
- [HUE-4403](#) - [aws] Close file after read using fast=True
- [HUE-4407](#) - [assist] Tables and views checkboxes can't be clicked
- [HUE-4410](#) - [editor] Fix editor size issue for large pasted queries
- [HUE-4413](#) - [dbms] Security: Full XSS in DBQuery editor
- [HUE-4414](#) - [editor] Allow to send queries even after closing the session
- [HUE-4415](#) - [editor] Disable "Save" button in query export unless output is filled
- [HUE-4416](#) - [editor] Do not save the coordinator constants in the json model
- [HUE-4417](#) - [aws] Switch from s3 to s3a to enable direct save to S3
- [HUE-4418](#) - [editor] Sessions tab doesn't display properly once the Settings is expanded
- [HUE-4419](#) - [editor] The new autocompleter should suggest columns from sub-queries
- [HUE-4421](#) - [editor] New saved query should be schedulable
- [HUE-4423](#) - [editor] Persist the coordinator ID of a submitted schedule
- [HUE-4424](#) - [librdbms] Support more than one database
- [HUE-4425](#) - [metastore] UX create table from a file, S3 or from a directory
- [HUE-4426](#) - [editor] The new autocompleter should work with all SQL dialects
- [HUE-4427](#) - [metastore] Filepicker should allow directories when selecting external table
- [HUE-4429](#) - [oozie] Generate unique workspace for managed workflows
- [HUE-4430](#) - [editor] Batch execution often fails to show any output
- [HUE-4431](#) - [editor] Batch mode does not always update the query
- [HUE-4432](#) - [editor] The new autocompleter should be more forgiving when editing incomplete UDFs in select list
- [HUE-4433](#) - [editor] The new autocompleter should handle analytical functions
- [HUE-4434](#) - [editor] Fixed header has a wrong offset in fullscreen when the editor is manually resized
- [HUE-4435](#) - [test] Add list_modules subcommand to `hue test`
- [HUE-4436](#) - [hive] Fix test_install_examples tests to point at the database configured in BeeswaxSampleProvide
- [HUE-4437](#) - [editor] The autocompleter should handle multiple db references in table primaries for Impala
- [HUE-4438](#) - [editor] Progress status and truncating warning when direct downloading results as Excel
- [HUE-4439](#) - [notebook] Fix notebook TestHiveserver2ApiWithHadoop.test_explain
- [HUE-4440](#) - [editor] Assist slightly scrolls horizontally
- [HUE-4441](#) - [editor] Set result menu as fixed
- [HUE-4442](#) - [editor] Improve column list scroll bars
- [HUE-4443](#) - [security] Automatically default secure_ssl_redirect to True when HTTPS is setup
- [HUE-4444](#) - [beeswax] test_analyze_table_and_read_statistics doesn't test correctly with Hive 2
- [HUE-4445](#) - [fb] Return correct result response on file upload
- [HUE-4448](#) - [hbase] Change test_list_tables hbase test to query each cluster name
- [HUE-4449](#) - [indexer] Add ruby log to indexer
- [HUE-4450](#) - [desktop] Fix HTTPS termination using X-Forwarded-Proto header
- [HUE-4451](#) - [editor] Add tooltips to result set legend
- [HUE-4452](#) - [metadata] Fix optimizer view name and update navigator API version to v9
- [HUE-4454](#) - [security] Disclosure of Web Server Information
- [HUE-4455](#) - [security] secure_content_security_policy breaks the editor and GA
- [HUE-4457](#) - [editor] Column legend horizontal scroll bar are missing until you scroll down
- [HUE-4459](#) - [indexer] Add morphline generation tests for each file format
- [HUE-4461](#) - [metastore] If create external table is selected, input path needs to be a directory
- [HUE-4462](#) - [editor] Status and progress report when downloading large Excel files
- [HUE-4463](#) - [editor] Result column search on column types too
- [HUE-4464](#) - [editor] Result column quick scroll does not scroll with a search

- [HUE-4465](#) - [notebook] Bump the darkness of the non selected icons
- [HUE-4466](#) - [security] Deliver csrftoken cookie with secure bit set if possible
- [HUE-4467](#) - [indexer] Add field operation morphline generation tests for each operation
- [HUE-4469](#) - [editor] Can have lot of white spaces below the query
- [HUE-4470](#) - [editor] Snippet properties for File filechooser always starts at /
- [HUE-4471](#) - [indexer] Add geoip operation automatically to hue format
- [HUE-4474](#) - [indexer] Let user set an existing field to unique id
- [HUE-4475](#) - [editor] Lines are not expanded when loading a long query from the cookie
- [HUE-4476](#) - [editor] Add RESET to the Hive highlight
- [HUE-4478](#) - [editor] Style the editor settings panel
- [HUE-4479](#) - [indexer] Allow indexing multiple files on the same collection
- [HUE-4480](#) - [meta] Filter SQL objects in the backend
- [HUE-4481](#) - [editor] Show a message when the Excel results have been truncated
- [HUE-4483](#) - [core] Do not use backticks in the ini files
- [HUE-4484](#) - [editor] Add column name and type to the title of column browser
- [HUE-4485](#) - [editor] Scroll to assist entries from search result
- [HUE-4486](#) - [editor] Add back sorting to the result table
- [HUE-4487](#) - [editor] Result table search
- [HUE-4488](#) - [editor] Show result row detail view
- [HUE-4489](#) - [editor] The new autocompleter should allow white space between UDF name and parenthesis
- [HUE-4490](#) - [oozie] Generate workflow from Java snippet
- [HUE-4491](#) - [fb] Enable creating buckets on S3
- [HUE-4492](#) - [indexer] Migrate to the oozie library to submit the java indexer job
- [HUE-4494](#) - [indexer] Add a load from JSON function to Create Index Wizard
- [HUE-4495](#) - [assist] Switching DB does not display the tables
- [HUE-4498](#) - [security] Fixed Content Security Policy blocks PDF in HBase app
- [HUE-4499](#) - [editor] Gradient map not always showing on the editor
- [HUE-4501](#) - [indexer] Add basic hive CSV file format outline to indexer
- [HUE-4502](#) - [editor] Fix issue with column type precision
- [HUE-4503](#) - [fb] Disable timestamp for S3 buckets and folders
- [HUE-4505](#) - [editor] Resultset search is too transparent
- [HUE-4506](#) - [editor] Provide more information why query couldn't be saved
- [HUE-4507](#) - [editor] Cannot see column names in the gradient map visualization
- [HUE-4508](#) - [editor] Search in result set is issuing a js error
- [HUE-4511](#) - [editor] Field scroll does not scroll or grey the column
- [HUE-4512](#) - [editor] Header and column not always aligned
- [HUE-4513](#) - [editor] Use weights for keyword autocompletion
- [HUE-4514](#) - [beeswax] Fix test_explain_query and test_explain_query_i18n for default execution engine Tez
- [HUE-4516](#) - [indexer] Switching file types doesn't update fields properly in the indexer wizard
- [HUE-4518](#) - [editor] Use different icon for the column list
- [HUE-4519](#) - [fb] Check for user permissions first before uploading
- [HUE-4521](#) - [editor] Hide/show column doesn't work on fixed rows
- [HUE-4522](#) - [beeswax] Remove max_rows limit for HiveServer2 get_configuration
- [HUE-4523](#) - [editor] Enable resizing of the column list panel
- [HUE-4524](#) - [hadoop] Fix test_yarn_ssl_validate which assumes default YARN_CLUSTER is defined
- [HUE-4526](#) - [oozie] API to generate sequential workflow with sequential document actions
- [HUE-4527](#) - [indexer] Date fields that aren't required can crash the indexing job
- [HUE-4530](#) - [indexer] Only warn when the collection name is empty
- [HUE-4531](#) - [fb] Enable deleting buckets on S3
- [HUE-4534](#) - [indexer] Fix syslog groking for smart indexer

- [HUE-4535](#) - [editor] Add autocomplete for CREATE and ALTER statements
- [HUE-4536](#) - [editor] Load query generates JS error
- [HUE-4538](#) - [search] Field detail popup cannot always be closed
- [HUE-4539](#) - [search] Remove field collapse icon
- [HUE-4540](#) - [job] Lower max number of jobs displayed on the main page
- [HUE-4542](#) - [editor] Column filter should be 100% of the column list width
- [HUE-4544](#) - [indexer] Extend Parquet morphline to support generic fields
- [HUE-4548](#) - [editor] Sample popup tab seems off
- [HUE-4549](#) - [core] Missing file for the collect static plugin
- [HUE-4550](#) - [fb] Fix file chooser for S3 after S3A migration
- [HUE-4551](#) - [editor] The new autocompleter should support CTEs (WITH) and UNION
- [HUE-4552](#) - [editor] Improve performance on checking/unchecking all fields
- [HUE-4553](#) - [oozie] The dashboard table scrolls horizontally
- [HUE-4554](#) - [editor] Status checking can go in infinite loop in batch mode
- [HUE-4555](#) - [core] Avoid opening file chooser on enter key
- [HUE-4556](#) - [core] Fix coverage failure due to missing source files
- [HUE-4559](#) - [indexer] Enable Select Fields button when a path has been entered into wizard
- [HUE-4560](#) - [core] Add permissions controls to authorize access to S3 across all components
- [HUE-4562](#) - [editor] Switching to marker map makes the page freeze
- [HUE-4563](#) - [impala] Handle case for non-string partition keys or values
- [HUE-4564](#) - [core] Log stderr on failure to coerce password from script
- [HUE-4565](#) - [core] Update natural sort algorithm
- [HUE-4566](#) - [metastore] Query created by Browse Data action should have a limit and semi-colon
- [HUE-4567](#) - [aws] Read AWS credentials from configured script files
- [HUE-4568](#) - [editor] Export a query to an index
- [HUE-4569](#) - [editor] If we don't have a query history select the saved query tab
- [HUE-4570](#) - [editor] Autocomplete popup disappears even when calling it when typing
- [HUE-4572](#) - [core] Fix HueDataTable with multiple instances on the same page
- [HUE-4573](#) - [editor] The autocompleter should support carriage return in statements
- [HUE-4574](#) - [editor] Enable autocompletion after '-' without whitespace
- [HUE-4575](#) - [core] Add notice to hue.ini regarding the custom maps and security
- [HUE-4576](#) - [editor] Fix autocompletion before OR in value expressions
- [HUE-4577](#) - [editor] The new autocompleter should support CREATE, DROP and SET ROLE
- [HUE-4579](#) - [editor] The new autocompleter should support Hive INSERT statements
- [HUE-4580](#) - [editor] The new autocompleter should support ANALYZE TABLE
- [HUE-4581](#) - [editor] The new autocompleter should support DROP FUNCTION
- [HUE-4582](#) - [editor] The new autocompleter should fully support LOAD
- [HUE-4584](#) - [editor] The new autocompleter should support DELETE
- [HUE-4586](#) - [editor] The new autocompleter should support GROUPING SETS, CUBE and ROLLUP
- [HUE-4587](#) - [editor] The new autocompleter should completely support EXPLAIN
- [HUE-4588](#) - [editor] The new autocompleter should support SET
- [HUE-4589](#) - [editor] The new autocompleter should completely support COMPUTE and DROP STATS
- [HUE-4590](#) - [editor] The new autocompleter should support INVALIDATE METADATA
- [HUE-4591](#) - [editor] The new autocompleter should support REFRESH
- [HUE-4598](#) - [editor] Empty table list sometimes the first time
- [HUE-4600](#) - [editor] Result table headers breaks on sort
- [HUE-4601](#) - [fb] Creating a bucket with a not allowed name on S3 will HTTP 500
- [HUE-4603](#) - [fb] Disable bucket move and copy
- [HUE-4604](#) - [editor] Locked rows does not go well with headers
- [HUE-4605](#) - [fb] Move and copy in S3 have the HDFS path autocompleters in the popup

- [HUE-4609](#) - [editor] Resizing the window should resize also the column list/grid
- [HUE-4610](#) - [editor] Avoid trembling of fixed headers on query editing
- [HUE-4611](#) - [editor] Export to S3 should not say HDFS
- [HUE-4614](#) - [editor] Query result search box always present when switching tabs
- [HUE-4618](#) - [editor] Query selection highlighting does not stay
- [HUE-4621](#) - [core] On loaddata failure, print stdout to logs
- [HUE-4623](#) - [editor] ON and ASC and DESC are not being highlighted on Impala
- [HUE-4624](#) - [editor] Doing enter in the query result search box could go to the next match
- [HUE-4625](#) - [fb] S3 filechooser popup do not display the parent directory as
- [HUE-4626](#) - [fb] S3 in the filechooser do not show up with FF
- [HUE-4628](#) - [useradmin] Improve and unify help icons
- [HUE-4631](#) - [home] DB transaction failing because of atomic block on home page
- [HUE-4632](#) - [oozie] Log which type of oozie action are created in workflows
- [HUE-4633](#) - [notebook] Disable by default Oozie integration until 3.12
- [HUE-4634](#) - [editor] The autocompleter should allow some errors in the select list
- [HUE-4637](#) - [editor] Scroll on column list is extremely slow on Linux
- [HUE-4638](#) - [editor] Infinite scroll blinks the fixed legend when fetching the next 100 batch
- [HUE-4640](#) - [editor] Fix for DB prefix autocomplete when typing table names after "? FROM db."
- [HUE-4641](#) - [indexer] GeoIP will stop indexing if an IP is missing from it's lookup table
- [HUE-4642](#) - [editor] Ace autocomplete should honour the weights for partial matching
- [HUE-4643](#) - [editor] Autocompleter is silent for the second argument of concat
- [HUE-4644](#) - [editor] Don't refresh the assist 5 seconds after loading a query
- [HUE-4645](#) - [core] Disable inline display of SVG files
- [HUE-4646](#) - [core] The initial welcome wizard redirects to the old home
- [HUE-4647](#) - [home] Better align new document dropdown
- [HUE-4648](#) - [security] The File ACLs tree doesn't render if there's a file in the HDFS root
- [HUE-4649](#) - [sqoop] SqoopResource got broken after change in Resource
- [HUE-4650](#) - [fb] Selecting S3A from filechooser is not displaying the S3 filesystem
- [HUE-4651](#) - [fb] Prevent inline display of non authorized mime types
- [HUE-4652](#) - [core] Add settings for django-axes AXES_BEHIND_REVERSE_PROXY and AXES_REVERSE_PROXY_HEADER
- [HUE-4657](#) - [indexer] A failed grok match attempt will stop indexing
- [HUE-4658](#) - [doc] Update release 3.11
- [HUE-4659](#) - [editor] The new autocompleter should merge columns from multiple tables when suggesting columns
- [HUE-4660](#) - [editor] Ace autocomplete spinner disappears too quick
- [HUE-4661](#) - [core] Demo backend should keep the admin user logged in
- [HUE-4665](#) - [editor] Autocompletion of lateral views should not have parenthesis around the column aliases
- [HUE-4666](#) - [search] Use the actual Hue user for impersonating in non secure mode
- [HUE-4667](#) - [editor] Geo plotting of states is case sensitive
- [HUE-4668](#) - [fb] S3 rename directory raises IOError
- [HUE-4670](#) - [metastore] Allow headers to be ignored/removed when creating external table
- [HUE-4671](#) - [fb] S3 create bucket should allow underscores in name
- [HUE-4673](#) - [fb] S3 bucket names should automatically lowercase them
- [HUE-4675](#) - [metastore] Assist not loaded and pointing to the source on the create table from a file page
- [HUE-4676](#) - [editor] Table header is not XSS safe
- [HUE-4678](#) - [editor] The label on marker maps is not updated to reflect what you choose on the left side dropdown
- [HUE-4682](#) - [editor] Scroll to a col, reexec the query, the greyed column is still there
- [HUE-4683](#) - [editor] Export result question mark icon does nothing
- [HUE-4684](#) - [editor] No placeholder showing up in new editors
- [HUE-4686](#) - [fb] Bubble up error when trying to create directory / in filechooser
- [HUE-4688](#) - [metastore] Authorize the selection of a directory in the create table from a file wizard

- [HUE-4689](#) - [fb] D&D a file into a directory in S3 gets 'RenameFormFormSet' object has no attribute 'data'
- [HUE-4691](#) - [fb] S3 D&D is allowed after a page refresh but not after opening a directory
- [HUE-4696](#) - [editor] JS .top error on certain reload/execute combos
- [HUE-4697](#) - [editor] The autocompleter should suggest select list aliases
- [HUE-4699](#) - [editor] Autocompleter fix for decimal values
- [HUE-4701](#) - [editor] Older saved queries throw "'NoneType' object has no attribute 'update_data'"
- [HUE-4703](#) - [Fixing[yarn]] Correct username is not used on hard failover
- [HUE-4704](#) - [security] Fixed Arbitrary host header accepted in Hue
- [HUE-4705](#) - [editor] Possible XSS when hovering on saved query description
- [HUE-4706](#) - [core] Importing documents should ignore reserved directories
- [HUE-4707](#) - [aws] Enable non-US region support for AWS S3
- [HUE-4708](#) - [aws] Add test_rename for directory rename in S3
- [HUE-4709](#) - [editor] Search does not load cells when navigating
- [HUE-4713](#) - [oozie] Spark example should use yarn client mode and show how to add dependencies
- [HUE-4718](#) - [core] jHueScrollUp pollutes the DOM with more than one scroll up anchor
- [HUE-4719](#) - [editor] Search disappears on load of new records
- [HUE-4720](#) - [oozie] Drag & Drop saved Spark app into a workflow
- [HUE-4724](#) - [editor] Closing the table search should un-highlight the selected cell
- [HUE-4725](#) - [editor] The autocompleter should handle Impala complex types in the table list
- [HUE-4727](#) - [fb] Raise warning if CHECKACCESS is not supported
- [HUE-4732](#) - [editor] Autocomplete should work inside parenthesized value expressions
- [HUE-4733](#) - [editor] Autocompleter doesn't support arithmetic operations in in value list
- [HUE-4734](#) - [editor] Autocompleter should trigger after '.' when autocomplete as you type is enabled
- [HUE-4735](#) - [editor] The autocompleter should also suggest tables with '.' when suggesting columns from multiple tables
- [HUE-4736](#) - [editor] The autocompleter throws JS error
- [HUE-4737](#) - [core] Upgrade boto to 2.42.0
- [HUE-4742](#) - [core] Data sample popup HueDataTable renders weirdly on FF
- [HUE-4747](#) - [editor] Download form should be submitted to a new tab otherwise the snippet gets closed
- [HUE-4753](#) - [editor] The autocompleter doesn't merge columns correctly
- [HUE-4754](#) - [editor] Increase default autocompleter API timeout to 5 seconds
- [HUE-4758](#) - [editor] The autocompleter should suggest keywords at the start of a statement
- [HUE-4759](#) - [editor] Don't suggest tables in the select list when there is a FROM clause
- [HUE-4760](#) - [core] Fix error with indexer config when search is disabled
- [HUE-4763](#) - [editor] Update the editor keyword highlight rules for Hive and Impala
- [HUE-4767](#) - [editor] Limit the autocompleter length before and after the cursor
- [HUE-4768](#) - [editor] Autocomplete is being called multiple times when the cursor is on top of '*' in a select statement
- [HUE-4769](#) - [editor] The autocompleter should support INSERT OVERWRITE DIRECTORY
- [HUE-4773](#) - [editor] Autocompleter spins forever when completing subqueries containing asterisk
- [HUE-4787](#) - [editor] Fixing Marker map tiles are not showing up
- [HUE-4801](#) - [core] Do not hide the username in small resolution
- [HUE-4804](#) - [search] Download function of HTML widget breaks the display
- [HUE-4809](#) - [oozie] Only add trustore paths when they are actually existing
- [HUE-4810](#) - [core] Fix tests by setting data to valid JSON type
- [HUE-4813](#) - [core] Disable collecting referrer URL in Google Analytics
- [HUE-4823](#) - [fb] Display exception when Hue is configured with invalid S3 access key
- [HUE-4827](#) - [fb] Pointing to a wrong S3 region breaks with no information
- [HUE-4871](#) - [useradmin] An unprivileged user can enumerate users
- [HUE-4882](#) - [editor] Redraw fixed header on column toggling
- [HUE-4891](#) - [useradmin] An unprivileged user can list document items

- [HUE-4910](#) - [editor] Improve result table scroll performance and fix header positioning
- [HUE-4916](#) - [core] Truncate last name to 30 chars on ldap import
- [HUE-4917](#) - [fb] S3 upload progress of large files is not consistent
- [HUE-4921](#) - [oozie] Workflow dashboard page should not fail when properties have unicode characters
- [HUE-4928](#) - [oozie] New Spark action can't be added in workflows
- [HUE-4935](#) - [aws] Enable support for AWS security token
- [HUE-4938](#) - [editor] Fix issue where the map crashes the browser for large coordinate values
- [HUE-4941](#) - [editor] Fixed Content Security Policy directive blocks an image when navigating on marker map
- [HUE-4968](#) - [oozie] Remove access to /oozie/import_workflow when v2 is enabled
- [HUE-4994](#) - [oozie] Consider default path for decision nodes in dashboard graph
- [HUE-4996](#) - [editor] Nested types can not be extended in assist
- [HUE-4997](#) - [impala] get_partitions is not being closed
- [HUE-4998](#) - [impala] get_configuration is not being closed
- [HUE-4999](#) - [impala] Set default Impala idle_session_timeout to 1 hour
- [HUE-5007](#) - [core] Get ride of bad utf8 character when doing make-locale
- [HUE-5007](#) - [core] Update localization for de, es, fr, ja, ko, zh
- [HUE-5014](#) - [core] Prevent copying a Document folder into itself which then creates a recursive depth issue
- [HUE-5040](#) - [notebook] Keep the progress bar to 99% orange until the status of query is READY
- [HUE-5041](#) - [editor] Hue export large file to HDFS doesn't work on non-default database
- [HUE-5050](#) - [core] Logout fails for local login when multiple backends are used
- [IMPALA-1112](#) - Remove some unnecessary code from cross-compilation
- [IMPALA-1240](#) - add back spilling sort now that sorter is not flaky
- [IMPALA-1440](#) - test for insert mem limit
- [IMPALA-3018](#) - Address various small memory allocation related bugs
- [IMPALA-1619](#) - Support 64-bit allocations
- [IMPALA-1633](#) - GetOperationStatus should set errorMessage and sqlState
- [IMPALA-1671](#) - Print time and link to coordinator web UI once query is submitted in shell
- [IMPALA-1683](#) - Allow REFRESH on a single partition
- [IMPALA-2347](#) - Reuse metastore client connections in Catalog
- [IMPALA-2459](#) - Implement next_day date/time UDF
- [IMPALA-2700](#) - ASCII NUL characters are doubled on insert into text tables
- [IMPALA-2767](#) - Web UI call to force expire sessions
- [IMPALA-2878](#) - Fix Base64Decode error and remove duplicate codes
- [IMPALA-2885](#) - ScannerContext::Stream objects should be owned by ScannerContext
- [IMPALA-2979](#) - Fix scheduling on remote hosts
- [IMPALA-3018](#) - Don't return NULL on zero length allocations
- [IMPALA-3063](#) - Separate join inversion from join ordering
- [IMPALA-3084](#) - Cache the sequence of table ref and materialized tuple ids during analysis
- [IMPALA-3181](#) - Add noexcept to some functions
- [IMPALA-3201](#) - buffer pool header only
- [IMPALA-3206](#) - Enable codegen for AVRO_DECIMAL
- [IMPALA-3210](#) - last/first_value() support for IGNORE NULLS
- [IMPALA-3223](#) - Remove boost multiprecision in thirdparty
- [IMPALA-3225](#) - Add script to push from gerrit to ASF
- [IMPALA-3227](#) - generate test TPC data sets during data load
- [IMPALA-3253](#) - Modify gen_build_version.sh to always output the right version
- [IMPALA-3336](#) - qgen: do not randomly generate query options
- [IMPALA-3376](#) - Extra definition level when writing Parquet files
- [IMPALA-3418](#) - The Impala FE project relies on Z-tools snapshot builds
- [IMPALA-3442](#) - Replace '> >' with '>>' in template decls

- [IMPALA-3449](#) - Kudu deploy.py should find clusters by displayName
- [IMPALA-3454](#) - Kudu deletes may fail if subqueries are used
- [IMPALA-3470](#) - DecompressorTest is flaky
- [IMPALA-3491](#) - Merge test_hbase_metadata.py into compute_stats.py. Use unique db fixture
- [IMPALA-3501](#) - ee tests: detect build type and support different timeouts based on the same
- [IMPALA-3507](#) - update binutils version to fix slow linking
- [IMPALA-3521](#) - Impalad should communicate with the statestore after binding to the hs2 and besswax ports
- [IMPALA-3530](#) - Clean up test_ddl.py. Part 1
- [IMPALA-3567](#) - Part 1: groundwork to make Join build sides DataSinks
- [IMPALA-3575](#) - Add retry to backend connection request and rpc timeout
- [IMPALA-3587](#) - Get rid of not_default_fs skip marker
- [IMPALA-3600](#) - Add missing admission control tests
- [IMPALA-3606](#) - Fix Java NPE when trying to add an existing partition
- [IMPALA-3611](#) - track unused Disk IO buffer memory
- [IMPALA-3627](#) - Clean up RPC structures in ImpalaInternalService
- [IMPALA-3632](#) - Add script for running cppclean over the BE code
- [IMPALA-3647](#) - track runtime filter memory in separate tracker
- [IMPALA-3656](#) - Hitting DCHECK/CHECK does not write minidumps
- [IMPALA-3664](#) - S3A test_keys_do_not_work fails
- [IMPALA-3674](#) - Lazy materialization of LLVM module bitcode
- [IMPALA-3677](#) - Write minidump on SIGUSR1
- [IMPALA-3682](#) - Don't retry unrecoverable socket creation errors
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3715](#) - Include total usage of JVM memory
- [IMPALA-3715](#) - Include more info by default in Impala debug memz webpage
- [IMPALA-3716](#) - Add Memory Tab in query's Details page
- [IMPALA-3727](#) - Change microbenchmarks to use percentile-based reporting
- [IMPALA-3729](#) - batch_size=1 coverage for avro scanner
- [IMPALA-3734](#) - C++11 - Replace boost:shared_ptr with std:: equivalent
- [IMPALA-3736](#) - Move Impala HTTP handlers to a separate class
- [IMPALA-3737](#) - Local filesystem build failed loading custom schemas
- [IMPALA-3751](#) - fix clang build errors and warnings
- [IMPALA-3753](#) - Disable create table test for old aggs and joins
- [IMPALA-3756](#) - Fix wrong argument type in HiveStringsTest
- [IMPALA-3757](#) - Add missing lock in RuntimeProfile::ComputeTimeInProfile
- [IMPALA-3762](#) - Download Python requirements before they are needed
- [IMPALA-3763](#) - download_requirements fixes
- [IMPALA-3764](#) - fuzz test HDFS scanners and fix parquet bugs found
- [IMPALA-3767](#) - bootstrap_virtualenv fails to find cython distribution
- [IMPALA-3774](#) - fix download_requirements for older Python versions
- [IMPALA-3778](#) - Fix ASF packaging build
- [IMPALA-3779](#) - Disable cache pool reader thread when HDFS isn't running
- [IMPALA-3780](#) - avoid many small reads past end of block
- [IMPALA-3786](#) - Remove "Cloudera" from impalad webpage title
- [IMPALA-3790](#) - AC tests timeout in codecoverage builds
- [IMPALA-3799](#) - Make MAX_SCAN_RANGE_LENGTH accept formatted quantities
- [IMPALA-3806](#) - remove a few modern shell idioms to improve RHEL5 support
- [IMPALA-3817](#) - Ensure filter hash function is the same on all hardware
- [IMPALA-3839](#) - Fix race condition in impala_cluster.py
- [IMPALA-3843](#) - Update warning for non-SSSE3 CPUs

- [IMPALA-3845](#) - Split up hdfs-parquet-scanner.cc into more files/components
- [IMPALA-3852](#) - Remove Derby and Shiro FE dependencies
- [IMPALA-3854](#) - Fix use-after-free in HdfsTextScanner::Close()
- [IMPALA-3856](#) - Fix BinaryPredicate normalization for Kudu
- [IMPALA-3857](#) - KuduScanNode race on returning "optional" threads
- [IMPALA-3864](#) - qgen: reduce likelihood of create_query() exceptions
- [IMPALA-3866](#) - consistent user-facing terminology for scratch dirs
- [IMPALA-3881](#) - Add DataTables 1.10.12 to www/
- [IMPALA-3886](#) - Improve log of pip_download.py
- [IMPALA-3892](#) - qgen: always run Impala with -convert_legacy_hive_parquet_utc_timestamps=true
- [IMPALA-3905](#) - Add HdfsScanner::GetNext() interface and implementation for Parquet
- [IMPALA-3906](#) - Materialize implicitly referenced IR functions
- [IMPALA-3914](#) - SKIP_TOOLCHAIN_BOOTSTRAP skips Python package downloads
- [IMPALA-3918](#) - Remove Cloudera copyrights and add ASF license header
- [IMPALA-3923](#) - fix overflow in BufferedTupleStream::GetRows()
- [IMPALA-3924](#) - Ubuntu16 support
- [IMPALA-3936](#) - BufferedBlockMgr fixes for Pin() while write in flight
- [IMPALA-3939](#) - Data loading may fail on tpch kudu
- [IMPALA-3943](#) - Do not throw scan errors for empty Parquet files
- [IMPALA-3946](#) - fix MemPool integrity issues with empty chunks
- [IMPALA-3952](#) - Clear scratch batch mem pool if Open() failed
- [IMPALA-3953](#) - Fixes for KuduScanNode BE test failure
- [IMPALA-3954](#) - Add unique_database to scanner test
- [IMPALA-3957](#) - Test failure in S3 build: TestLoadData.test_load
- [IMPALA-3964](#) - Fix crash when a count(*) is performed on a nested collection
- [IMPALA-3969](#) - stress test: add option to set common query options
- [IMPALA-3972](#) - Improve display of /varz page
- [IMPALA-3992](#) - bad shell error message when running nonexistent file
- [IMPALA-3223](#) - Fix the include path for Hadoop's header files
- [IMPALA-3253](#) - Modify gen_build_version.sh to always output the right version
- [IMPALA-3832](#) - Merge "fix DCHECK calling AddDetail() on an OK status" into cdh5-trunk
- [IMPALA-3832](#) - fix DCHECK calling AddDetail() on an OK status
- [IMPALA-3905](#) - Fixes for the modified HdfsScanner interface
- [OOZIE-1814](#) - Oozie should mask any passwords in logs and REST interfaces
- [OOZIE-2244](#) - Oozie should mask passwords in the logs when logging command arguments
- [OOZIE-2349](#) - Method getCoordJobInfo(String jobId, String filter, int offset, int length, boolean desc) is not present in LocalOozieClientCoord
- [OOZIE-2362](#) - SQL injection in BulkJPAExecutor
- [OOZIE-2402](#) - oozie-setup.sh sharelib create takes a long time on large clusters
- [OOZIE-2447](#) - Illegal character 0x0 oozie client
- [OOZIE-2462](#) - When calling ParamChecker.notNull() in CoordActionsIgnoreXCommand.java, 'Action' should be passed instead of 'Action cannot be null'
- [OOZIE-2467](#) - Oozie can shutdown itself on long GC pause
- [OOZIE-2475](#) - Oozie does not cleanup action dir of killed actions
- [OOZIE-2476](#) - When one of the action from fork fails with transient error, WF never joins
- [OOZIE-2493](#) - TestDistcpMain deletes action.xml from wrong filesystem
- [OOZIE-2515](#) - Duplicate information for 'Changing endtime/pausetime of a Bundle Job' in CommandLineTool wiki
- [OOZIE-2516](#) - Update web service documentation for jobs API
- [OOZIE-2539](#) - Incorrect property key is used for 'hive log4j configuration file for execution mode'
- [OOZIE-2541](#) - Possible resource leak in Hive2Credentials

- [OOZIE-2542](#) - Option to disable OpenJPA BrokerImpl finalization
- [OOZIE-2546](#) - Improperly closed resources in OozieDBCLI
- [OOZIE-2551](#) - Feature request: epoch timestamp generation
- [OOZIE-2552](#) - Update ActiveMQ version for security and other fixes
- [OOZIE-2553](#) - Cred tag is required for all actions in the workflow even if an action does not require it
- [OOZIE-2556](#) - TestAbandonedCoordChecker.testCatchupJob is flaky
- [OOZIE-2574](#) - Oozie to support replication-enabled mysql urls
- [OOZIE-2578](#) - Oozie example distcp job fails to run within an encrypted zone with checksum match error
- [OOZIE-2579](#) - Bulk kill tests in TestBulkWorkflowXCommand might fail because of a race condition
- [OOZIE-2581](#) - Oozie should reset SecurityManager in finally block
- [OOZIE-2587](#) - Disable SchedulerService on certain tests
- [OOZIE-2603](#) - Give thread pools a meaningful name in CallableQueueService and SchedulerService
TestCoordActionsKillXCommand.testActionKillCommandActionNumbers and testActionKillCommandDate
- [OOZIE-2623](#) - Oozie should use a dummy OutputFormat
- [OOZIE-2632](#) - Provide database dump/load functionality to make database migration easier
- [OOZIE-2660](#) - Create documentation for DB Dump/Load functionality
- [PIG-2949](#) - JsonLoader only reads arrays of objects
- [PIG-3413](#) - JsonLoader fails the pig job in case of malformed json input
- [PIG-3619](#) - Update piggybank to include XPath function
- [PIG-3664](#) - Piggy Bank XPath UDF can't be called
- [PIG-3730](#) - Performance issue in SelfSpillBag
- [PIG-3970](#) - Increase PermGen size, tests ran out of memory
- [PIG-4355](#) - Piggybank: XPath cant handle namespace in xpath, nor can it return more than one match
- [PIG-4787](#) - CHD:14987: BackportLog JSONLoader exception while parsing records
- [SENTRY-1201](#) - Sentry ignores database prefix for MSCK statement
- [SENTRY-1208](#) - Make HOST implied in privileges if not specified explicitly
- [SENTRY-1228](#) - SimpleFileProviderBackend error message, remove missing spaces
- [SENTRY-1230](#) - Add filesystem tests to test Sentry with user data storage on S3
- [SENTRY-1231](#) - Sentry doesn't secure index location uri, when do "CREATE INDEX LOCATION "/uri"
- [SENTRY-1252](#) - grantServerPrivilege and revokeServerPrivilege should treat * and ALL as synonyms when action is not explicitly specified
- [SENTRY-1253](#) - SentryShellKafka incorrectly sets component as "KAFKA"
- [SENTRY-1265](#) - Sentry service should not require a TGT
- [SENTRY-1269](#) - Converter vs Convertor - spelling inconsistent
- [SENTRY-1292](#) - Reorder DBModelAction EnumSet
- [SENTRY-1293](#) - Avoid converting string permission to Privilege object
- [SENTRY-1299](#) - Add a test case to verify SentryStore#verifySentryStoreSchema works
- [SENTRY-1311](#) - Improve usability of URI privileges by supporting mixed use of URIs with and without scheme
- [SENTRY-1320](#) - Truncate table db_name.table_name fails
- [SENTRY-1334](#) - backport[column level privileges] test and add test for CTAS and Create View AS SELECT (cross databases cases)
- [SENTRY-1345](#) - ACLS on table folder disappear after insert for unpartitioned tables
- [SENTRY-1376](#) - Fix alter property case correctly - Deletes ACLS on the table
- [SENTRY-1378](#) - Login fails for a secure Sentry Web UI
- [SENTRY-1416](#) - kafka-sentry tool service name default is different from KafkaSentryAuthorizer default service name
- [SENTRY-1429](#) - TestHDFSIntegration improvements
- [SENTRY-1454](#): Fix intermittent timeout issue for TestHDFSIntegration
- [SENTRY-1447](#) - When s3 is configured as HDFS defaultFS and Hive Warehouse Dir, need to fix some e2e test failures
- [SENTRY-1450](#) - Have privilege converter set by Kafka binding

- [SENTRY-1453](#) - Enable passing sentry client cache configs from Kafka conf
- [SENTRY-1233](#) - Logging improvements for `SentryConfigToolSolr`.
- [SENTRY-1119](#) - Allow data engines to obtain the `ActionFactory` directly from the configuration, instead of having hardcoded component-specific classes. This will allow external data engines to integrate with Sentry easily.
- [SENTRY-1229](#) - Added a basic configurable cache to `SentryGenericProviderBackend`.
- [SOLR-5750](#) - Backup/Restore API for SolrCloud
- [SOLR-5922](#) - Add support for adding core properties to SolrJ Collection Admin Request calls
- [SOLR-6637](#) - Solr should have a way to restore a core
- [SOLR-6761](#) - Ability to ignore commit and optimize requests from clients when running in SolrCloud mode
- [SOLR-7294](#) - Migrate API fails with 'Invalid status request: notfoundretried 6times' message
- [SOLR-7374](#) - Backup/Restore should provide a param for specifying the directory implementation it should use
- [SOLR-7766](#) - Support creation of a coreless collection
- [SOLR-8449](#) - Multiple restores on the same core does not work
- [SOLR-9055](#) - Make collection backup/restore extensible
- [SOLR-9242](#) - Collection level backup/restore should provide a param for specifying the repository implementation it should use
- [SOLR-9269](#) - Ability to create/delete/list snapshots for a solr core
- [SOLR-9310](#) - PeerSync fails on a node restart due to `IndexFingerPrint` mismatch
- [SOLR-9326](#) - Ability to create/delete/list snapshots for a solr collection
- [SOLR-9441](#) - Solr collection backup on HDFS can only be manipulated by the Solr process owner
- [SPARK-10372](#) - [CORE] Basic test framework for entire spark scheduler
- [SPARK-12009](#) - [YARN] Avoid reallocating YARN container when driver wants to stop all executors
- [SPARK-12941](#) - [SQL][MASTER] Spark-SQL JDBC Oracle dialect fails to map string datatypes to Oracle VARCHAR datatype
- [SPARK-13904](#) - [SCHEDULER] Add support for pluggable cluster manager
- [SPARK-14881](#) - [PYTHON] [SPARKR] pyspark and sparkR shell default log level should match spark-shell/Scala
- [SPARK-15714](#) - [CORE] Fix flaky o.a.s.scheduler.BlacklistIntegrationSuite
- [SPARK-15754](#) - [YARN] Not letting the credentials containing HDFS delegation tokens to be added in current user credential
- [SPARK-15865](#) - [CORE] Blacklist should not result in job hanging with less than 4 executors
- [SPARK-16711](#) - `YarnShuffleService` doesn't re-init properly on YARN rolling upgrade
- [SPARK-17171](#) - [WEB UI] DAG will list all partitions in the graph
- [SPARK-17433](#) - `YarnShuffleService` doesn't handle moving credentials levelDb
- [SPARK-17611](#) - [YARN][TEST] Make shuffle service test really test authentication
- [SPARK-17644](#) - [CORE] Do not add failedStages when abortStage for fetch failure
- [SPARK-17696](#) - [CORE] Partial backport of to branch-1.6
- [SQOOP-2938](#) - Mainframe import module extension to support data sets on tape
- [SQOOP-2952](#) - Fixing bug
- [SQOOP-3021](#) - `ClassWriter` fails if a column name contains a backslash character
- [ZOOKEEPER-2405](#) - `getTGT()` in `Login.java` mishandles confidential information
- [ZOOKEEPER-2477](#) - Documentation should refer to Java CLI shell and not C CLI shell

Issues Fixed in CDH 5.8.x

The following topics describe issues fixed in CDH 5.8.x, from newest to oldest release. You can also review [What's New In CDH 5.8.x](#) on page 1827 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.8.5

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.8.5:

- [CRUNCH-592](#) - Job fails for null `ByteBuffer` value in Avro tables.

- [FLUME-1899](#) - Make SpoolDir work with subdirectories
- [FLUME-2171](#) - Add Interceptor to remove headers from event
- [FLUME-2652](#) - Documented transaction handling semantics incorrect in developer guide.
- [FLUME-2797](#) - Use SourceCounter for SyslogTcpSource
- [FLUME-2798](#) - Malformed Syslog messages can lead to OutOfMemoryException
- [FLUME-2812](#) - Fix semaphore leak causing java.lang.Error: Maximum permit count exceeded in MemoryChannel
- [FLUME-2844](#) - SpillableMemoryChannel must start ChannelCounter
- [FLUME-2889](#) - Fixes to DateTime computations
- [FLUME-2901](#) - Document Kerberos setup for Kafka channel
- [FLUME-2910](#) - AsyncHBaseSink: Failure callbacks should log the exception that caused them
- [FLUME-2913](#) - Don't strip SLF4J from imported classpaths
- [FLUME-2918](#) - Speed up TaildirSource on directories with many files
- [FLUME-2922](#) - Sync SequenceFile.Writer before calling hflush
- [FLUME-2923](#) - Bump asynchbase version to 1.7.0
- [FLUME-2934](#) - Document new cachePatternMatching option for TaildirSource
- [FLUME-2935](#) - Bump java target version to 1.7
- [FLUME-2948](#) - docs: Fix parameters on Replicating Channel Selector example
- [FLUME-2954](#) - Make raw data appearing in log messages explicit
- [FLUME-2963](#) - FlumeUserGuide: Fix error in Kafka Source properties table
- [FLUME-2972](#) - Handle offset migration in the new Kafka Channel
- [FLUME-2975](#) - docs: Fix NetcatSource example
- [FLUME-2982](#) - Add localhost escape sequence to HDFS sink
- [FLUME-2983](#) - Handle offset migration in the new Kafka Source
- [FLUME-2999](#) - Kafka channel and sink should enable statically assigned partition per event via header
- [FLUME-3020](#) - Improve HDFS Sink escape sequence substitution
- [FLUME-3027](#) - Change Kafka Channel to clear offsets map after commit
- [FLUME-3031](#) - Change sequence source to reset its counter for event body on channel exception
- [FLUME-3049](#) - Make HDFS sink rotate more reliably in secure mode
- [HADOOP-7930](#) - Kerberos relogin interval in UserGroupInformation should be configurable
- [HADOOP-8436](#) - NPE In getLocalPathForWrite (path, conf) when the required context item is not configured
- [HADOOP-8437](#) - getLocalPathForWrite should throw IOException for invalid paths
- [HADOOP-8934](#) - Shell command ls should include sort options
- [HADOOP-10048](#) - LocalDirAllocator should avoid holding locks while accessing the filesystem
- [HADOOP-10300](#) - Allowed deferred sending of call responses.
- [HADOOP-10971](#) - Add -C flag to make `hadoop fs -ls` print filenames only
- [HADOOP-11031](#) - Design Document for Credential Provider API
- [HADOOP-11361](#) - Fix a race condition in MetricsSourceAdapter.updateJmxCache
- [HADOOP-11400](#) - GraphiteSink does not reconnect to Graphite after 'broken pipe'
- [HADOOP-11469](#) - KMS should skip default.key.acl and whitelist.key.acl when loading key acl.
- [HADOOP-11599](#) - Client#getTimeout should use IPC_CLIENT_PING_DEFAULT when IPC_CLIENT_PING_KEY is not configured
- [HADOOP-11619](#) - FTPFileSystem should override getDefaultPort.
- [HADOOP-11901](#) - BytesWritable fails to support 2G chunks due to integer overflow
- [HADOOP-12252](#) - LocalDirAllocator should not throw NPE with empty string configuration
- [HADOOP-12453](#) - Support decoding KMS Delegation Token with its own Identifier
- [HADOOP-12483](#) - Maintain wrapped SASL ordering for postponed IPC responses.
- [HADOOP-12537](#) - S3A to support Amazon STS temporary credentials
- [HADOOP-12548](#) - Read s3a creds from a Credential Provider
- [HADOOP-12609](#) - Fix intermittent failure of TestDecayRpcScheduler.
- [HADOOP-12655](#) - TestHttpServer.testBindAddress bind port range is wider than expected.

- [HADOOP-12659](#) - Incorrect usage of config parameters in token manager of KMS
- [HADOOP-12672](#) - RPC timeout should not override IPC ping interval
- [HADOOP-12723](#) - S3A: Add ability to plug in any AWSCredentialsProvider
- [HADOOP-12963](#) - Allow using path style addressing for accessing the s3 endpoint.
- [HADOOP-12973](#) - Make DU pluggable.
- [HADOOP-12974](#) - Create a CachingGetSpaceUsed implementation that uses df
- [HADOOP-12975](#) - Add jitter to CachingGetSpaceUsed's thread
- [HADOOP-13034](#) - Log message about input options in distcp lacks some items
- [HADOOP-13072](#) - WindowsGetSpaceUsed constructor should be public
- [HADOOP-13079](#) - Add -q option to Ls to print ? instead of non-printable characters
- [HADOOP-13132](#) - Handle ClassCastException on AuthenticationException in LoadBalancingKMSClientProvider
- [HADOOP-13155](#) - Implement TokenRenewer to renew and cancel delegation tokens in KMS
- [HADOOP-13189](#) - FairCallQueue makes callQueue larger than the configured capacity
- [HADOOP-13251](#) - Authenticate with Kerberos credentials when renewing KMS delegation token
- [HADOOP-13255](#) - KMSClientProvider should check and renew tgt when doing delegation token operations
- [HADOOP-13263](#) - Reload cached groups in background after expiry.
- [HADOOP-13270](#) - BZip2CompressionInputStream finds the same compression marker twice in corner case, causing duplicate data blocks
- [HADOOP-13317](#) - Add logs to KMS server-side to improve supportability
- [HADOOP-13353](#) - LdapGroupsMapping getPassword shouldn't return null when IOException throws
- [HADOOP-13381](#) - KMS clients should use KMS Delegation Tokens from current UGI
- [HADOOP-13433](#) - Race in UGI.reloginFromKeytab
- [HADOOP-13434](#) - Add bash quoting to Shell class.
- [HADOOP-13437](#) - KMS should reload whitelist and default key ACLs when hot-reloading
- [HADOOP-13457](#) - Remove hardcoded absolute path for shell executable.
- [HADOOP-13487](#) - Hadoop KMS should load old delegation tokens from Zookeeper on startup
- [HADOOP-13503](#) - Improve SaslRpcClient failure logging
- [HADOOP-13526](#) - Add detailed logging in KMS for the authentication failure of proxy user
- [HADOOP-13558](#) - UserGroupInformation created from a Subject incorrectly tries to renew the Kerberos ticket
- [HADOOP-13579](#) - Fix source-level compatibility after HADOOP-11252
- [HADOOP-13590](#) - Retry until TGT expires even if the UGI renewal thread encountered exception.
- [HADOOP-13627](#) - Have an explicit KerberosAuthException for UGI to throw, text from public constants
- [HADOOP-13638](#) - KMS should set UGI's Configuration object properly
- [HADOOP-13641](#) - Update UGI#spawnAutoRenewalThreadForUserCreds to reduce indentation
- [HADOOP-13669](#) - Addendum patch 2 for KMS Server should log exceptions before throwing.
- [HADOOP-13693](#) - Remove the message about HTTP OPTIONS in SPNEGO initialization message from kms audit log.
- [HADOOP-13749](#) - KMSClientProvider combined with KeyProviderCache can result in wrong UGI being used
- [HADOOP-13805](#) - UGI.getCurrentUser() fails if user does not have a keytab associated
- [HADOOP-13838](#) - KMSTokenRenewer should close providers
- [HADOOP-13953](#) - Make FTPFileSystem's data connection mode and transfer mode configurable
- [HADOOP-14003](#) - Make additional KMS tomcat settings configurable
- [HADOOP-14104](#) - Client should always ask namenode for kms provider path
- [HADOOP-14195](#) - CredentialProviderFactory\$getProviders is not thread-safe
- [HDFS-4176](#) - EditLogTailer should call rollEdits with a timeout.
- [HDFS-4210](#) - Throw helpful exception when DNS entry for JournalNode cannot be resolved
- [HDFS-6434](#) - Default permission for creating file should be 644 for WebHdfs/HttpFS
- [HDFS-6962](#) - ACLs inheritance conflict with umaskmode
- [HDFS-7413](#) - Some unit tests should use NameNodeProtocols instead of FSNameSystem
- [HDFS-7415](#) - Move FSNameSystem.resolvePath() to FSDirectory

- [HDFS-7420](#) - Delegate permission checks to FSDirectory
- [HDFS-7463](#) - Simplify FSNamesystem#getBlockLocationsUpdateTimes
- [HDFS-7478](#) - Move org.apache.hadoop.hdfs.server.namenode.NNConf to FSNamesystem
- [HDFS-7517](#) - Remove redundant non-null checks in FSNamesystem#getBlockLocations
- [HDFS-7597](#) - DelegationTokenIdentifier should cache the TokenIdentifier to UGI mapping
- [HDFS-7964](#) - Add support for async edit logging
- [HDFS-8224](#) - Schedule a block for scanning if its metadata file is corrupt
- [HDFS-8269](#) - getBlockLocations() does not resolve the .reserved path and generates incorrect edit logs when updating the atime
- [HDFS-8581](#) - ContentSummary on / skips further counts on yielding lock
- [HDFS-8709](#) - Clarify automatic sync in FSEditLog#logEdit.
- [HDFS-8829](#) - Make SO_RCVBUF and SO_SNDBUF size configurable for DataTransferProtocol sockets and allow configuring auto-tuning
- [HDFS-8897](#) - Balancer should handle fs.defaultFS trailing slash in HA
- [HDFS-9038](#) - DFS reserved space is erroneously counted towards non-DFS used.
- [HDFS-9085](#) - Show renewer information in DelegationTokenIdentifier#toString
- [HDFS-9137](#) - DeadLock between DataNode#refreshVolumes and BPOfferService#registrationSucceeded.
- [HDFS-9141](#) - Thread leak in Datanode#refreshVolumes.
- [HDFS-9259](#) - Make SO_SNDBUF size configurable at DFSClient side for hdfs write scenario.
- [HDFS-9276](#) - Failed to Update HDFS Delegation Token for long running application in HA mode
- [HDFS-9365](#) - Balancer does not work with the HDFS-6376 HA setup.
- [HDFS-9461](#) - DiskBalancer: Add Report Command
- [HDFS-9530](#) - ReservedSpace is not cleared for abandoned Blocks
- [HDFS-9601](#) - NNThroughputBenchmark.BlockReportStats should handle NotReplicatedYetException on adding block.
- [HDFS-9630](#) - DistCp minor refactoring and clean up
- [HDFS-9638](#) - Improve DistCp Help and documentation.
- [HDFS-9700](#) - DFSClient and DFSOutputStream should set TCP_NODELAY on sockets for DataTransferProtocol
- [HDFS-9732](#) - Improve DelegationTokenIdentifier.toString() for better logging
- [HDFS-9781](#) - FsDatasetImpl#getBlockReports can occasionally throw NullPointerException
- [HDFS-9805](#) - Add server-side configuration for enabling TCP_NODELAY for DataTransferProtocol and default it to true
- [HDFS-9820](#) - Improve distcp to support efficient restore to an earlier snapshot
- [HDFS-9906](#) - Remove spammy log spew when a datanode is restarted.
- [HDFS-9939](#) - Increase DecompressorStream skip buffer size
- [HDFS-9958](#) - BlockManager#createLocatedBlocks can throw NPE for corruptBlocks on failed storages
- [HDFS-10216](#) - Distcp -diff throws exception when handling relative path
- [HDFS-10270](#) - TestJMXGet:testNameNode() fails
- [HDFS-10298](#) - Document the usage of distcp -diff option
- [HDFS-10312](#) - Large block reports may fail to decode at NameNode due to 64 MB protobuf maximum length restriction
- [HDFS-10313](#) - Distcp need to enforce the order of snapshot names passed to -diff.
- [HDFS-10336](#) - TestBalancer failing intermittently because of not resetting UserGroupInformation completely
- [HDFS-10381](#) - DataStreamer DataNode exclusion log message should be warning.
- [HDFS-10397](#) - Distcp should ignore -delete option if -diff option is provided instead of exiting
- [HDFS-10403](#) - DiskBalancer: Add cancel command
- [HDFS-10457](#) - DataNode should not auto-format block pool directory if VERSION is missing.
- [HDFS-10481](#) - HTTPFS server should correctly impersonate as end user to open file
- [HDFS-10500](#) - Diskbalancer: Print out information when a plan is not generated
- [HDFS-10501](#) - DiskBalancer: Use the default datanode port if port is not provided

- [HDFS-10512](#) - VolumeScanner may terminate due to NPE in DataNode.reportBadBlocks
- [HDFS-10516](#) - Fix bug when warming up EDEK cache of more than one encryption zone
- [HDFS-10517](#) - DiskBalancer: Support help command
- [HDFS-10525](#) - Fix NPE in CacheReplicationMonitor#rescanCachedBlockMap
- [HDFS-10541](#) - Diskbalancer: When no actions in plan, error message says "Plan was generated more than 24 hours ago"
- [HDFS-10544](#) - Balancer doesn't work with IPFailoverProxyProvider.
- [HDFS-10552](#) - DiskBalancer "-query" results in NPE if no plan for the node
- [HDFS-10556](#) - DistCpOptions should be validated automatically
- [HDFS-10559](#) - DiskBalancer: Use SHA1 for Plan ID
- [HDFS-10567](#) - Improve plan command help message
- [HDFS-10588](#) - False alarm in datanode log - ERROR - Disk Balancer is not enabled
- [HDFS-10598](#) - DiskBalancer does not execute multi-steps plan
- [HDFS-10600](#) - PlanCommand#getThrsoldPercentage should not use throughput value.
- [HDFS-10609](#) - Uncaught InvalidEncryptionKeyException during pipeline recovery may abort downstream applications
- [HDFS-10641](#) - TestBlockManager#testBlockReportQueueing fails intermittently.
- [HDFS-10643](#) - Namenode should use loginUser(hdfs) to generateEncryptedKey
- [HDFS-10681](#) - DiskBalancer: query command should report Plan file path apart from PlanID.
- [HDFS-10715](#) - NPE when applying AvailableSpaceBlockPlacementPolicy
- [HDFS-10722](#) - Fix race condition in TestEditLog#testBatchedSyncWithClosedLogs
- [HDFS-10760](#) - DataXceiver#run() should not log InvalidToken exception as an error
- [HDFS-10763](#) - Open files can leak permanently due to inconsistent lease update
- [HDFS-10822](#) - Log DataNodes in the write pipeline.
- [HDFS-10879](#) - TestEncryptionZonesWithKMS#testReadWrite fails intermittently
- [HDFS-10963](#) - Reduce log level when network topology cannot find enough datanodes
- [HDFS-11012](#) - Unnecessary INFO logging on DFSclients for InvalidToken
- [HDFS-11040](#) - Add documentation for HDFS-9820 distcp improvement
- [HDFS-11056](#) - Concurrent append and read operations lead to checksum error
- [HDFS-11160](#) - VolumeScanner reports write-in-progress replicas as corrupt incorrectly
- [HDFS-11229](#) - HDFS-11056 failed to close meta file
- [HDFS-11275](#) - Check groupEntryIndex and throw a helpful exception on failures when removing ACL.
- [HDFS-11292](#) - log lastWrittenTxId etc info in logSyncAll
- [HDFS-11306](#) - Print remaining edit logs from buffer if edit log can't be rolled
- [HDFS-11363](#) - Need more diagnosis info when seeing Slow waitForAckedSeqno.
- [HDFS-11379](#) - DFSInputStream may infinite loop requesting block locations
- [HDFS-11689](#) - New exception thrown by DFSClient%isHDFSEncryptionEnabled broke hacky hive code
- [MAPREDUCE-4784](#) - TestRecovery occasionally fails
- [MAPREDUCE-6172](#) - TestDbClasses timeouts are too aggressive
- [MAPREDUCE-6359](#) - In RM HA setup, Cluster tab links populated with AM hostname instead of RM
- [MAPREDUCE-6442](#) - Stack trace is missing when error occurs in client protocol provider's constructor
- [MAPREDUCE-6473](#) - Job submission can take a long time during Cluster initialization
- [MAPREDUCE-6571](#) - JobEndNotification info logs are missing in AM container syslog
- [MAPREDUCE-6628](#) - Potential memory leak in CryptoOutputStream
- [MAPREDUCE-6633](#) - AM should retry map attempts if the reduce task encounters compression related errors
- [MAPREDUCE-6641](#) - TestTaskAttempt fails in trunk
- [MAPREDUCE-6670](#) - TestJobListCache#testEviction sometimes fails on Windows with timeout
- [MAPREDUCE-6680](#) - JHS UserLogDir scan algorithm sometime could skip directory with update in CloudFS (Azure FileSystem, S3, etc
- [MAPREDUCE-6718](#) - add progress log to JHS during startup
- [MAPREDUCE-6728](#) - Give fetchers hint when ShuffleHandler rejects a shuffling connection

- [MAPREDUCE-6738](#) - TestJobListCache.testAddExisting failed intermittently in slow VM testbed
- [MAPREDUCE-6761](#) - Regression when handling providers - invalid configuration ServiceConfiguration causes Cluster initialization failure
- [MAPREDUCE-6763](#) - Shuffle server listen queue is too small
- [MAPREDUCE-6771](#) - RMContainerAllocator sends container diagnostics event after corresponding completion event
- [MAPREDUCE-6798](#) - Fix intermittent failure of TestJobHistoryParsing.testJobHistoryMethods
- [MAPREDUCE-6817](#) - The format of job start time in JHS is different from those of submit and finish time.
- [MAPREDUCE-6839](#) - TestRecovery.testCrashed failed
- [YARN-2306](#) - Add test for leakage of reservation metrics in fair scheduler..
- [YARN-2336](#) - Fair scheduler's REST API returns a missing '[' bracket JSON for deep queue tree
- [YARN-2605](#) - [RM HA] Rest api endpoints doing redirect incorrectly.
- [YARN-2977](#) - Fixed intermittent TestNMClient failure.
- [YARN-3251](#) - Fixed a deadlock in CapacityScheduler when computing absoluteMaxAvailableCapacity in LeafQueue
- [YARN-3601](#) - Fix UT TestRMFailover.testRMWebAppRedirect
- [YARN-3654](#) - ContainerLogsPage web UI should not have meta-refresh
- [YARN-3722](#) - Merge multiple TestWebAppUtils into o.a.h.yarn.webapp.util.TestWebAppUtils
- [YARN-3933](#) - FairScheduler: Multiple calls to completedContainer are not safe.
- [YARN-3957](#) - FairScheduler NPE In FairSchedulerQueueInfo causing scheduler page to return 500.
- [YARN-4004](#) - container-executor should print output of docker logs if the docker container exits with non-0 exit status
- [YARN-4017](#) - container-executor overuses PATH_MAX
- [YARN-4092](#) - Addendum. Fixed UI redirection to print useful messages when both RMs are in standby mode
- [YARN-4245](#) - Generalize config file handling in container-executor
- [YARN-4255](#) - container-executor does not clean up docker operation command files
- [YARN-4363](#) - In TestFairScheduler, testcase should not create FairScheduler redundantly.
- [YARN-4411](#) - RMAAppAttemptImpl#createApplicationAttemptReport throws IllegalArgumentException
- [YARN-4459](#) - container-executor should only kill process groups
- [YARN-4544](#) - All the log messages about rolling monitoring interval are shown with WARN level
- [YARN-4555](#) - TestDefaultContainerExecutor#testContainerLaunchError fails on non-english locale environment
- [YARN-4556](#) - TestFifoScheduler.testResourceOverCommit fails
- [YARN-4820](#) - ResourceManager web redirects in HA mode drops query parameters
- [YARN-4866](#) - FairScheduler: AMs can consume all vcores leading to a livelock when using FAIR policy.
- [YARN-4878](#) - Expose scheduling policy and max running apps over JMX for Yarn queues.
- [YARN-4940](#) - yarn node -list -all failed if RM start with decommissioned node
- [YARN-4989](#) - TestWorkPreservingRMRestart#testCapacitySchedulerRecovery fails intermittently
- [YARN-5001](#) - Aggregated Logs root directory is created with wrong group if nonexistent
- [YARN-5048](#) - DelegationTokenRenewer#skipTokenRenewal may throw NPE
- [YARN-5077](#) - Fix FSLeafQueue#getFairShare() for queues with zero fairshare.
- [YARN-5107](#) - TestContainerMetrics fails.
- [YARN-5136](#) - Error in handling event type APP_ATTEMPT_REMOVED to the scheduler
- [YARN-5246](#) - NMWebAppFilter web redirects drop query parameters
- [YARN-5272](#) - Handle queue names consistently in FairScheduler.
- [YARN-5608](#) - TestAMRMClient.setup() fails with ArrayOutOfBoundsException
- [YARN-5704](#) - Provide config knobs to control enabling/disabling new/work in progress features in container-executor
- [YARN-5752](#) - TestLocalResourcesTrackerImpl#testLocalResourceCache times out
- [YARN-5837](#) - NPE when getting node status of a decommissioned node after an RM restart
- [YARN-5859](#) - TestResourceLocalizationService#testParallelDownloadAttemptsForPublicResource sometimes fails
- [YARN-5862](#) - TestDiskFailures.testLocalDirsFailures failed
- [YARN-5890](#) - FairScheduler should log information about AM-resource-usage and max-AM-share for queues

- [YARN-5920](#) - Fix deadlock in TestRMHA.testTransitionedToStandbyShouldNotHang
- [YARN-6042](#) - Dump scheduler and queue state information into FairScheduler DEBUG log.
- [YARN-6042](#) - Revert "Dump scheduler and queue state information into FairScheduler DEBUG log."
- [YARN-6042](#) - Dump scheduler and queue state information into FairScheduler DEBUG log.
- [YARN-6151](#) - FS preemption does not consider child queues over fairshare if the parent is under.
- [YARN-6175](#) - FairScheduler: Negative vcore for resource needed to preempt.
- [YARN-6264](#) - AM not launched when a single vcore is available on the cluster.
- [YARN-6359](#) - TestRM#testApplicationKillAtAcceptedState fails rarely due to race condition
- [YARN-6360](#) - Prevent FS state dump logger from cramming other log files
- [YARN-6453](#) - fairscheduler-statedump.log gets generated regardless of service
- [HBASE-12949](#) - Scanner can be stuck in infinite loop if the HFile is corrupted
- [HBASE-14644](#) - Region in transition metric is broken
- [HBASE-14818](#) - user_permission does not list namespace permissions
- [HBASE-14963](#) - Remove use of Guava Stopwatch from HBase client code
- [HBASE-15125](#) - BackportHBaseFsck's adoptHdfsOrphan function creates region with wrong end key boundary
- [HBASE-15324](#) - Jitter may cause desiredMaxFileSize overflow in ConstantSizeRegionSplitPolicy and trigger unexpected split
- [HBASE-15328](#) - sanity check the redirect used to send master info requests to the embedded regionserver.
- [HBASE-15378](#) - Scanner cannot handle heartbeat message with no results
- [HBASE-15430](#) - Failed taking snapshot - Manifest proto-message too large
- [HBASE-15465](#) - userPermission returned by getUserPermission() for the selected namespace does not have namespace set
- [HBASE-15496](#) - Throw RowTooBigException only for user scan/get
- [HBASE-15587](#) - FSTableDescriptors.getDescriptor() logs stack trace erroneously
- [HBASE-15613](#) - TestNamespaceCommand times out
- [HBASE-15621](#) - Suppress Hbase SnapshotHFile cleaner error messages when a snapshot is going on
- [HBASE-15683](#) - Min latency in latency histograms are emitted as Long.MAX_VALUE
- [HBASE-15698](#) - Increment TimeRange not serialized to server
- [HBASE-15746](#) - Remove extra RegionCoprocessor preClose() in RSRpcServices#closeRegion
- [HBASE-15808](#) - Reduce potential bulk load intermediate space usage and waste
- [HBASE-15856](#) - Don't cache unresolved addresses for connections
- [HBASE-15873](#) - ACL for snapshot restore / clone is not enforced
- [HBASE-15925](#) - provide default values for hadoop compat module related properties that match default hadoop profile.
- [HBASE-15931](#) - Add log for long-running tasks in AsyncProcess
- [HBASE-15955](#) - Disable action in CatalogJanitor#setEnabled should wait for active cleanup scan to finish
- [HBASE-16032](#) - Possible memory leak in StoreScanner
- [HBASE-16056](#) - Procedure v2 - fix master crash for FileNotFound
- [HBASE-16062](#) - Improper error handling in WAL Reader/Writer creation
- [HBASE-16093](#) - Fix splits failed before creating daughter regions leave meta inconsistent
- [HBASE-16135](#) - PeerClusterZnode under rs of removed peer may never be deleted
- [HBASE-16146](#) - Counters are expensive...
- [HBASE-16172](#) - Unify the retry logic in ScannerCallableWithReplicas and RpcRetryingCallerWithReadReplicas
- [HBASE-16194](#) - Should count in MSLAB chunk allocation into heap size change when adding duplicate cells
- [HBASE-16195](#) - Should not add chunk into chunkQueue if not using chunk pool in HeapMemStoreLAB
- [HBASE-16207](#) - can't restore snapshot without "Admin" permission
- [HBASE-16227](#) - [Shell] Column value formatter not working in scans. Tested : manually using shell.
- [HBASE-16238](#) - It's useless to catch SESSIONEXPIRED exception and retry in RecoverableZooKeeper
- [HBASE-16270](#) - Handle duplicate clearing of snapshot in region replicas
- [HBASE-16284](#) - Unauthorized client can shutdown the cluster

- [HBASE-16288](#) - HFile intermediate block level indexes might recurse forever creating multi TB files
- [HBASE-16294](#) - hbck reporting "No HDFS region dir found" for replicas
- [HBASE-16304](#) - HRegion#RegionScannerImpl#handleFileNotFoundException may lead to deadlock when trying to obtain write lock on updatesLock
- [HBASE-16317](#) - revert all ESAPI changes
- [HBASE-16319](#) - Fix TestCacheOnWrite after HBASE-16288
- [HBASE-16321](#) - ensure no findbugs-jsr305
- [HBASE-16340](#) - exclude Xerces implementation jars from coming in transitively.
- [HBASE-16345](#) - RpcRetryingCallerWithReadReplicas#call() should catch some RegionServer Exceptions
- [HBASE-16350](#) - Undo server abort from HBASE-14968
- [HBASE-16360](#) - TableMapReduceUtil addHBaseDependencyJars has the wrong class name for PrefixTreeCodec
- [HBASE-16429](#) - FSHLog: deadlock if rollWriter called when ring buffer filled with appends
- [HBASE-16460](#) - Can't rebuild the BucketAllocator's data structures when BucketCache uses FileIOEngine
- [HBASE-16604](#) - Scanner retries on IOException can cause the scans to miss data
- [HBASE-16662](#) - Fix open POODLE vulnerabilities
- [HBASE-16699](#) - Overflows in AverageIntervalRateLimiter's refill() and getWaitInterval()
- [HBASE-16721](#) - Concurrency issue in WAL unflushed seqld tracking
- [HBASE-16767](#) - Mob compaction needs to clean up files in /hbase/mobdir/.tmp and /hbase/mobdir/.tmp/.bulkload when running into IO exceptions
- [HBASE-16807](#) - , RegionServer will fail to report new active Hmaster until HMaster/RegionServer failover.
- [HBASE-16824](#) - Writer.flush() can be called on already closed streams in WAL roll
- [HBASE-16841](#) - Data loss in MOB files after cloning a snapshot and deleting that snapshot
- [HBASE-16931](#) - Setting cell's seqld to zero in compaction flow might cause RS down.
- [HBASE-16960](#) - RegionServer hang when aborting
- [HBASE-17020](#) - keylen in midkey() dont computed correctly
- [HBASE-17023](#) - Region left unassigned due to AM and SSH each thinking others would do the assignment work
- [HBASE-17044](#) - Fix merge failed before creating merged region leaves meta inconsistent
- [HBASE-17058](#) - Lower epsilon used for jitter verification from HBASE-15324
- [HBASE-17069](#) - RegionServer writes invalid META entries for split daughters in some circumstances
- [HBASE-17072](#) - CPU usage starts to climb up to 90-100% when using G1GC; purge ThreadLocal usage
- [HBASE-17206](#) - FSHLog may roll a new writer successfully with unflushed entries
- [HBASE-17241](#) - Avoid compacting already compacted mob files with _del files
- [HBASE-17265](#) - Region left unassigned in master failover when region failed to open
- [HBASE-17275](#) - Assign timeout may cause region to be unassigned forever
- [HBASE-17328](#) - Properly dispose of looped replication peers
- [HBASE-17381](#) - ReplicationSourceWorkerThread can die due to unhandled exceptions
- [HBASE-17409](#) - Limit jsonp callback name to prevent xss
- [HBASE-17452](#) - Failed taking snapshot - region Manifest proto-message too large
- [HBASE-17522](#) - Handle JVM throwing runtime exceptions when we ask for details on heap usage the same as a correctly returned 'undefined'.
- [HBASE-17558](#) - ZK dumping jsp should escape HTML.
- [HBASE-17561](#) - table status page should escape values that may contain arbitrary characters.
- [HBASE-17675](#) - ReplicationEndpoint should choose new sinks if a SaslException occurs
- [HBASE-17717](#) - Explicitly use "sasl" ACL scheme for hbase superuser
- [HIVE-6758](#) - Beeline doesn't work with -e option when started in background
- [HIVE-7443](#) - Fix HiveConnection to communicate with Kerberized Hive JDBC server and alternative JDKs
- [HIVE-7723](#) - Explain plan for complex query with lots of partitions is slow due to in-efficient collection used to find a matching ReadEntity
- [HIVE-10007](#) - Support qualified table name in analyze table compute statistics for columns
- [HIVE-10384](#) - BackportRetryingMetaStoreClient does not retry wrapped TTransportExceptions

- [HIVE-10728](#) - deprecate `unix_timestamp(void)` and make it deterministic
- [HIVE-10965](#) - direct SQL for stats fails in 0-column case
- [HIVE-11028](#) - Tez: table self join and join with another table fails with `IndexOutOfBoundsException`
- [HIVE-11141](#) - Improve `RuleRegExp` when the Expression node stack gets huge
- [HIVE-11243](#) - Changing log level in `Utilities.getBaseWork`
- [HIVE-11375](#) - Broken processing of queries containing NOT (`x IS NOT NULL` and `x <> 0`)
- [HIVE-11428](#) - Performance: Struct `IN()` clauses are extremely slow
- [HIVE-11432](#) - Hive macro give same result for different arguments
- [HIVE-11487](#) - Add `getNumPartitionsByFilter` api in metastore api
- [HIVE-11594](#) - Analyze Table for column names with embedded spaces
- [HIVE-11671](#) - Optimize `RuleRegExp` in DPP codepath
- [HIVE-11717](#) - nohup mode is not support for new hive
- [HIVE-11747](#) - Unnecessary error log is shown when executing a "INSERT OVERWRITE LOCAL DIRECTORY" cmd in the embedded mode
- [HIVE-11827](#) - STORED AS AVRO fails SELECT COUNT(*) when empty
- [HIVE-11842](#) - Improve `RuleRegExp` by caching some internal data structures
- [HIVE-11849](#) - NPE in `HiveHBaseTableShapshotInputFormat` in query with just count(*)
- [HIVE-11901](#) - `StorageBasedAuthorizationProvider` requires write permission on table for SELECT statements
- [HIVE-11980](#) - Follow up on HIVE-11696, exception is thrown from CTAS from the table with table-level serde is Parquet while partition-level serde is JSON
- [HIVE-12077](#) - MSCK Repair table should fix partitions in batches
- [HIVE-12083](#) - HIVE-10965 introduces thrift error if `partNames` or `colNames` are empty
- [HIVE-12179](#) - Add option to not add `spark-assembly.jar` to Hive classpath
- [HIVE-12277](#) - Hive macro results on `macro_duplicate.q` different after adding ORDER BY
- [HIVE-12349](#) - NPE in ORC SARG for IS NULL queries on Timestamp and Date columns
- [HIVE-12465](#) - Hive might produce wrong results when (outer) joins are merged
- [HIVE-12475](#) - Parquet schema evolution within `array<struct<>>` doesn't work
- [HIVE-12556](#) - Ctrl-C in beeline doesn't kill Tez query on HS2
- [HIVE-12619](#) - Switching the field order within an array of structs causes the query to fail
- [HIVE-12635](#) - Hive should return the latest hbase cell timestamp as the row timestamp value
- [HIVE-12768](#) - Thread safety: binary sortable serde decimal deserialization
- [HIVE-12780](#) - Fix the output of the history command in Beeline
- [HIVE-12785](#) - View with union type and UDF to the struct is broken
- [HIVE-12834](#) - Fix to accept the arrow keys in BeeLine CLI
- [HIVE-12891](#) - Hive fails when `java.io.tmpdir` is set to a relative location
- [HIVE-12976](#) - `MetaStoreDirectSql` doesn't batch IN lists in all cases
- [HIVE-13043](#) - Reload function has no impact to function registry
- [HIVE-13058](#) - Add session and operation_log directory deletion messages
- [HIVE-13090](#) - Hive metastore crashes on NPE with `ZooKeeperTokenStore`
- [HIVE-13129](#) - `ClService` leaks HMS connection
- [HIVE-13149](#) - Remove some unnecessary HMS connections from HS2
- [HIVE-13198](#) - Authorization issues with cascading views
- [HIVE-13237](#) - Select parquet struct field with upper case throws NPE
- [HIVE-13240](#) - `GroupByOperator`: Drop the hash aggregates when closing operator
- [HIVE-13372](#) - Hive Macro overwritten when multiple macros are used in one column
- [HIVE-13381](#) - Timestamp & date should have precedence in type hierarchy than string group
- [HIVE-13429](#) - Tool to remove dangling scratch dir
- [HIVE-13462](#) - `HiveResultSetMetaData.getPrecision()` fails for NULL columns
- [HIVE-13539](#) - `HiveHFileOutputFormat` searching the wrong directory for HFiles
- [HIVE-13590](#) - Kerberized HS2 with LDAP auth enabled fails in multi-domain LDAP case

- [HIVE-13620](#) - Merge llap branch work to master
- [HIVE-13625](#) - Hive Prepared Statement when executed with escape characters in parameter fails
- [HIVE-13645](#) - Beeline needs null-guard around hiveVars and hiveConfVars read
- [HIVE-13704](#) - Don't call DistCp.execute() instead of DistCp.run()
- [HIVE-13736](#) - View's input/output formats are TEXT by default.
- [HIVE-13749](#) - Memory leak in Hive Metastore
- [HIVE-13864](#) - Beeline ignores the command that follows a semicolon and comment
- [HIVE-13866](#) - flatten callstack for directSQL errors
- [HIVE-13884](#) - Disallow queries in HMS fetching more than a configured number of partitions
- [HIVE-13895](#) - HoS start-up overhead in yarn-client mode
- [HIVE-13932](#) - Hive SMB Map Join with small set of LIMIT failed with NPE
- [HIVE-13936](#) - Add streaming support for row_number
- [HIVE-13953](#) - Issues in HiveLockObject equals method
- [HIVE-13991](#) - Union All on view fail with no valid permission on underneath table
- [HIVE-13997](#) - Insert overwrite directory doesn't overwrite existing files
- [HIVE-14006](#) - Hive query with UNION ALL fails with ArrayIndexOutOfBoundsException.
- [HIVE-14015](#) - SMB MapJoin failed for Hive on Spark when kerberized
- [HIVE-14037](#) - java.lang. ClassNotFoundException for the jar in hive.reloadable.aux.jars.path in mapreduce
- [HIVE-14055](#) - directSql - getting the number of partitions is broken
- [HIVE-14098](#) - Logging task properties, and environment variables might contain passwords
- [HIVE-14118](#) - Make the alter partition exception more meaningful
- [HIVE-14137](#) - Hive on Spark throws FileAlreadyExistsException for jobs with multiple empty tables
- [HIVE-14142](#) - java.lang. ClassNotFoundException for the jar in hive.reloadable.aux.jars.path for Hive on Spark
- [HIVE-14173](#) - NPE was thrown after enabling directsql in the middle of session
- [HIVE-14187](#) - JDOPersistenceManager objects remain cached if MetaStoreClient#close is not called
- [HIVE-14198](#) - Refactor aux jar related code to make them more consistent
- [HIVE-14209](#) - Add some logging info for session and operation management
- [HIVE-14210](#) - ExecDriver should call jobclient.close() to trigger cleanup
- [HIVE-14296](#) - Session count is not decremented when HS2 clients do not shutdown cleanly.
- [HIVE-14342](#) - Beeline output is garbled when executed from a remote shell
- [HIVE-14383](#) - SparkClientImpl should pass principal and keytab to spark-submit instead of calling kinit explicitly
- [HIVE-14421](#) - FS.deleteOnExit holds references to _tmp_space.db files.
- [HIVE-14229](#) - The jars in hive. aux.jar.paths are not added to session classpath
- [HIVE-14436](#) - Hive 1.2.1/Hitting "ql.Driver: FAILED: IllegalArgumentException Error: , expected at the end of 'decimal(9'" after enabling hive.optimize.skewjoin and with MR engine
- [HIVE-14457](#) - Partitions in encryption zone are still trashed though an exception is returned
- [HIVE-14519](#) - Multi insert query bug
- [HIVE-14538](#) - beeline throws exceptions with parsing hive config when using !sh statement
- [HIVE-14693](#) - Some partitions will be left out when partition number is the multiple of the option hive.msck.repair.batch.size
- [HIVE-14697](#) - Can not access kerberized HS2 Web UI
- [HIVE-14715](#) - Hive throws NumberFormatException with query with Null value
- [HIVE-14743](#) - ArrayIndexOutOfBoundsException - HBASE-backed views' query with JOINS
- [HIVE-14762](#) - Add logging while removing scratch space
- [HIVE-14764](#) - Enabling "hive.metastore.metrics.enabled" throws OOM in HiveMetastore
- [HIVE-14784](#) - Operation logs are disabled automatically if the parent directory does not exist.
- [HIVE-14799](#) - Query operation are not thread safe during its cancellation
- [HIVE-14805](#) - Subquery inside a view will have the object in the subquery as the direct input
- [HIVE-14817](#) - Shutdown the SessionManager timeoutChecker thread properly upon shutdown.
- [HIVE-14819](#) - FunctionInfo for permanent functions shows TEMPORARY FunctionType

- [HIVE-14820](#) - RPC server for spark inside HS2 is not getting server address properly
- [HIVE-15054](#) - Hive insertion query execution fails on Hive on Spark
- [HIVE-15061](#) - Metastore types are sometimes case sensitive
- [HIVE-15090](#) - Temporary DB failure can stop ExpiredTokenRemover thread
- [HIVE-15231](#) - query on view with CTE and alias fails with table not found error
- [HIVE-15291](#) - Comparison of timestamp fails if only date part is provided
- [HIVE-15338](#) - Wrong result from non-vectorized DATEDIFF with scalar parameter of type DATE/TIMESTAMP
- [HIVE-15346](#) - "values temp table" should not be an input
- [HIVE-15410](#) - WebHCat supports get/set table property with its name containing period and hyphen
- [HIVE-15517](#) - NOT (x <=> y) returns NULL if x or y is NULL
- [HIVE-15551](#) - memory leak in directsql for mysql+bonecp specific initialization
- [HIVE-15572](#) - Improve the response time for query canceling when it happens during acquiring locks
- [HIVE-15735](#) - In some cases, view objects inside a view do not have parents.
- [HIVE-15782](#) - query on parquet table returns incorrect result when hive.optimize.index.filter is set to true
- [HIVE-15872](#) - The PERCENTILE_APPROX UDAF does not work with empty set
- [HIVE-15997](#) - Resource leaks when query is cancelled
- [HIVE-16047](#) - Shouldn't try to get KeyProvider unless encryption is enabled
- [HIVE-16156](#) - FileSinkOperator should delete existing output target when renaming
- [HIVE-16175](#) - Possible race condition in InstanceCache
- [HIVE-16394](#) - HoS does not support queue name change in middle of session
- [HUE-3065](#) - [oozie] Sub-workflow submitted from coordinator gets parent workflow graph
- [HUE-3079](#) - [oozie] Some links of a Fork can point to deleted nodes
- [HUE-4147](#) - [useradmin] Ignore (objectclass=*) filter when searching for LDAP users
- [HUE-4386](#) - [oozie] Remove oozie.coord.application.path from properties when rerunning workflow
- [HUE-4462](#) - [oozie] Fix deployment_dir for the bundle in oozie example fixtures
- [HUE-4706](#) - [useradmin] update AuthenticationForm to allow activated users to login
- [HUE-4921](#) - [core] Skip idle session timeout relogin popup on running job jobs call when idle session timeout is disabled
- [HUE-4941](#) - [jobbrowser] Unable to kill jobs with Resource Manager HA enabled
- [HUE-4969](#) - [security] Can't type any / in the HDFS ACLs path input
- [HUE-5158](#) - [editor] Older queries after upgrade do not provide direct save
- [HUE-5390](#) - [editor] Improve import testing of beeswax queries to notebook format
- [HUE-5659](#) - [home] Ignore history dependencies when importing document from different cluster
- [HUE-5679](#) - [yarn] Reset API_CACHE on logout
- [HUE-5714](#) - [yarn] Fix unittest for MR API Cache
- [HUE-6090](#) - [search] Typing in the search bar always redirect to the end of the input
- [HUE-6131](#) - [metastore] No information surfaced when LOAD data from Create table from file fails
- [HUE-6133](#) - [editor] Horizontal scrollbar can be hidden under the first fixed column
- [HUE-6144](#) - [editor] Make it possible to turn autocomplete on or off
- [HUE-6197](#) - [editor] Enable scrolling past the end of the editor
- [HUE-6228](#) - [editor] API for progress status and truncating warning when direct downloading results as Excel
- [IMPALA-1346](#) - /1590/2344: fix sorter buffer mgmt when spilling
- [IMPALA-1619](#), [IMPALA-3018](#): Address various small memory allocation related bugs
- [IMPALA-1619](#) - Support 64-bit allocations.
- [IMPALA-1657](#) - Rework detection and reporting of corrupt table stats.
- [IMPALA-2864](#) - Ensure that client connections are closed after a failed Open()
- [IMPALA-3018](#) - Don't return NULL on zero length allocations.
- [IMPALA-3159](#) - impala-shell does not accept wildcard or SAN certificates
- [IMPALA-3167](#) - Fix assignment of WHERE conjunct through grouping agg + OJ.
- [IMPALA-3314](#) - Fix Avro schema loading for partitioned tables.

- [IMPALA-3344](#) - Simplify sorter and document/enforce invariants.
- [IMPALA-3441](#), [IMPALA-3659](#): check for malformed Avro data
- [IMPALA-3499](#) - Split catalog update
- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3575](#) - Add retry to backend connection request and rpc timeout
- [IMPALA-3628](#) - Fix cancellation from shell when security is enabled
- [IMPALA-3633](#) - cancel fragment if coordinator is gone
- [IMPALA-3646](#) - Handle corrupt RLE literal or repeat counts of 0.
- [IMPALA-3670](#) - fix sorter buffer mgmt bugs
- [IMPALA-3678](#) - Fix migration of predicates into union operands with an order by + limit.
- [IMPALA-3680](#) - Cleanup the scan range state after failed hdfs cache reads
- [IMPALA-3682](#) - Don't retry unrecoverable socket creation errors
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3711](#) - Remove unnecessary privilege checks in getDbsMetadata()
- [IMPALA-3732](#) - handle string length overflow in avro files
- [IMPALA-3745](#) - parquet invalid data handling
- [IMPALA-3751](#) - fix clang build errors and warnings
- [IMPALA-3754](#) - fix TestParquet.test_corrupt_rle_counts flakiness
- [IMPALA-3776](#) - fix 'describe formatted' for Avro tables
- [IMPALA-3820](#) - Handle linkage errors while loading Java UDFs in Catalog
- [IMPALA-3861](#) - Replace BetweenPredicates with their equivalent CompoundPredicate.
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3884](#) - Support TYPE_TIMESTAMP for HashTableCtx::CodegenAssignNullValue()
- [IMPALA-3915](#) - Register privilege and audit requests when analyzing resolved table refs.
- [IMPALA-3930](#), [IMPALA-2570](#): Fix shuffle insert hint with constant partition exprs.
- [IMPALA-3940](#) - Fix getting column stats through views.
- [IMPALA-3949](#) - Log the error message in FileSystemUtil.copyToLocal()
- [IMPALA-3964](#) - Fix crash when a count(*) is performed on a nested collection.
- [IMPALA-3965](#) - TSSLSocketWithWildcardSAN.py not exported as part of impala-shell build lib
- [IMPALA-3983](#), [IMPALA-3974](#): Delete function jar resources after load
- [IMPALA-4019](#) - initialize member variables in HdfsTableSink
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4037](#), [IMPALA-4038](#): fix locking during query cancellation
- [IMPALA-4049](#) - fix empty batch handling NLJ build side
- [IMPALA-4076](#) - Fix runtime filter sort compare method
- [IMPALA-4099](#) - Fix the error message while loading UDFs with no JARs
- [IMPALA-4120](#) - Incorrect results with LEAD() analytic function
- [IMPALA-4135](#) - Thrift threaded server times-out connections during high load
- [IMPALA-4153](#) - Fix count(*) on all blank('') columns - test
- [IMPALA-4170](#) - Fix identifier quoting in COMPUTE INCREMENTAL STATS.
- [IMPALA-4180](#) - Synchronize accesses to RuntimeState::reader_contexts_
- [IMPALA-4196](#) - Cross compile bit-byte-functions
- [IMPALA-4223](#) - Handle truncated file read from HDFS cache
- [IMPALA-4237](#) - Fix materialization of 4 byte decimals in data source scan node.
- [IMPALA-4246](#) - SleepForMs() utility function has undefined behavior for > 1s
- [IMPALA-4260](#) - Alter table add column drops all the column stats
- [IMPALA-4263](#) - Fix wrong omission of agg/analytic hash exchanges.
- [IMPALA-4266](#) - Java udf returning string can give incorrect results
- [IMPALA-4282](#) - Remove max length check for type strings.
- [IMPALA-4293](#) - query profile should include error log

- [IMPALA-4295](#) - XFAIL wildcard SSL test
- [IMPALA-4336](#) - Cast exprs after unnesting union operands.
- [IMPALA-4363](#) - Add Parquet timestamp validation
- [IMPALA-4383](#) - Ensure plan fragment report thread is always started
- [IMPALA-4391](#) - fix dropped statuses in scanners
- [IMPALA-4423](#) - Correct but conservative implementation of Subquery.equals().
- [IMPALA-4433](#) - Always generate testdata using the same time zone setting
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog This commit fixes an issue where multiple long-running operations on the same catalog object (e.g. table) can block other catalog operations from making progress.
- [IMPALA-4488](#) - HS2 GetOperationStatus() should keep session alive
- [IMPALA-4518](#) - CopyStringVal() doesn't copy null string
- [IMPALA-4539](#) - fix bug when scratch batch references I/O buffers
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots
- [IMPALA-4579](#) - SHOW CREATE VIEW fails for view containing a subquery
- [IMPALA-4765](#) - Avoid using several loading threads on one table.
- [IMPALA-4767](#) - Workaround for HIVE-15653 to preserve table stats.
- [IMPALA-4779](#), [IMPALA-4780](#): Fix conditional functions built-in and Timestamp bounds
- [IMPALA-4787](#) - Optimize APPX_MEDIAN() memory usage
- [IMPALA-4916](#) - Fix maintenance of set of item sets in DisjointSet.
- [IMPALA-4995](#) - Fix integer overflow in TopNNNode::PrepareForOutput
- [IMPALA-4997](#) - Fix overflows in Sorter::Tupliterator
- [IMPALA-5005](#) - Don't allow server to send SASL COMPLETE msg out of order
- [IMPALA-5088](#) - Fix heap buffer overflow
- [IMPALA-5253](#) - Use appropriate transport for StatestoreSubscriber
- [IMPALA-4391](#) - fix dropped status in scanners
- [OOZIE-1814](#) - Oozie should mask any passwords in logs and REST interfaces
- [OOZIE-2068](#) - Configuration as part of sharelib
- [OOZIE-2194](#) - oozie job -kill doesn't work with spark action
- [OOZIE-2243](#) - Kill Command does not kill the child job for java action
- [OOZIE-2314](#) - Unable to kill old instance child job by workflow or coord rerun by Launcher
- [OOZIE-2329](#) - Make handling yarn restarts configurable
- [OOZIE-2345](#) - Parallel job submission for forked actions
- [OOZIE-2347](#) - Remove unnecessary new Configuration()/new jobConf() calls from oozie
- [OOZIE-2436](#) - Fork/join workflow fails with oozie.action.yarn.tag must not be null
- [OOZIE-2504](#) - Create a log4j.properties under HADOOP_CONF_DIR in Shell Action
- [OOZIE-2533](#) - Patch-1550 - workaround for
- [OOZIE-2555](#) - Oozie SSL enable setup does not return port for admin -servers
- [OOZIE-2567](#) - HCat connection is not closed while getting hcat cred
- [OOZIE-2584](#) - Eliminate Thread.sleep() calls in TestMemoryLocks
- [OOZIE-2589](#) - CompletedActionXCommand is hardcoded to wrong priority
- [OOZIE-2649](#) - Can't override sub-workflow configuration property if defined in parent workflow XML
- [OOZIE-2656](#) - OozieShareLibCLI uses op system username instead of Kerberos to upload jars
- [OOZIE-2678](#) - Oozie job -kill doesn't work with tez jobs
- [OOZIE-2739](#) - Remove property expansion pattern from ShellMain's log4j properties content
- [OOZIE-2742](#) - Unable to kill applications based on tag
- [OOZIE-2777](#) - Config-default.xml longer than 64k results in java.io.UTFDataFormatException
- [OOZIE-2818](#) - Can't overwrite oozie.action.max.output.data on a per-workflow basis
- [PIG-3807](#) - Pig creates wrong schema after dereferencing nested tuple fields with sorts
- [PIG-3818](#) - PIG-2499 is accidentally reverted
- [PIG-3970](#) - Increase PermGen size, tests ran out of memory

- [PIG-4052](#) - TestJobControlSleep, TestInvokerSpeed are unreliable
- [SENTRY-1201](#) - Sentry ignores database prefix for MSCK statement
- [SENTRY-1265](#) - Sentry service should not require a TGT as it is not talking to other kerberos services as a client
- [SENTRY-1311](#) - Improve usability of URI privileges by supporting mixed use of URIs with and without scheme
- [SENTRY-1313](#) - Database prefix is not honoured when executing grant statement
- [SENTRY-1345](#) - ACLs on table folder disappear after insert for unpartitioned tables
- [SENTRY-1520](#) - Provide mechanism for triggering HMS full snapshot
- [SOLR-5776](#) - backportEnabled SSL tests can easily exhaust random generator entropy and block. Set the server side to SHA1PRNG as in Steve's original patch. Use less SSL in a test run. refactor SSLConfig so that SSLTestConfig can provide SSLContexts using a NullSecureRandom to prevent SSL tests from blocking on entropy starved machines Alternate (psuedo random) NullSecureRandom for Constants.SUN_OS replace NullSecureRandom w/ NotSecurePsuedoRandom
- [SOLR-6295](#) - Fix child filter query creation to never match parent docs in SolrExampleTests
- [SOLR-7280](#) - /Missing test resources
- [SOLR-7280](#) - BackportLoad cores in sorted order and tweak coreLoadThread counts to improve cluster stability on restarts
- [SOLR-7866](#) - Harden code to prevent an unhandled NPE when trying to determine the max value of the version field.
- [SOLR-9091](#) - ZkController#publishAndWaitForDownStates logic is inefficient
- [SOLR-9236](#) - AutoAddReplicas will append an extra /tlog to the update log location on replica failover.
- [SOLR-9284](#) - The HDFS BlockDirectoryCache should not let it's keysToRelease or names maps grow indefinitely.
- [SOLR-9310](#), [SOLR-9524](#)
- [SOLR-9330](#) - Fix AlreadyClosedException on admin/mbeans?stats=true
- [SOLR-9699](#), [SOLR-4668](#): fix exception from core status in parallel with core reload
- [SOLR-9819](#) - Upgrade Apache commons-fileupload to 1.3.2, fixing a security vulnerability
- [SOLR-9848](#) - Lower solr.cloud.wait-for-updates-with-stale-state-pause back down from 7 seconds.
- [SOLR-9859](#) - backport ofreplication.properties cannot be updated after being written and neither replication.properties or index.properties are durable in the face of a crash. Don't log error on NoSuchFileException
- [SOLR-9901](#) - backport of SOLR-9899 Implement move in HdfsDirectoryFactory. SOLR-9899: StandardDirectoryFactory should use optimizations for all FilterDirectorys not just NRTCachingDirectory.
- [SOLR-10031](#) - Validation of filename params in ReplicationHandler
- [SOLR-10114](#) - backport of SOLR-9941 - Reordered delete-by-query can delete or omit child documents
- [SOLR-10119](#) - TestReplicationHandler assertion fixes part of
- [SOLR-10121](#), [SOLR-10116](#): BlockCache corruption with high concurrency
- [SOLR-10338](#) - backportConfigure SecureRandom non blocking for tests.
- [SPARK-8428](#) - [SPARK-13850] Fix integer overflows in TimSort
- [SPARK-12009](#) - [YARN] Avoid to re-allocating yarn container while driver want to stop all Executors
- [SPARK-12241](#) - [YARN] Improve failure reporting in Yarn client obtainTokenForHBase()
- [SPARK-12339](#) - [SPARK-11206][WEBUI] Added a null check that was removed in
- [SPARK-12392](#) - [CORE] Optimize a location order of broadcast blocks by considering preferred local hosts
- [SPARK-12523](#) - [YARN] Support long-running of the Spark On HBase and hive meta store.
- [SPARK-12941](#) - [SQL][MASTER] Spark-SQL JDBC Oracle dialect fails to map string datatypes to Oracle VARCHAR datatype mapping
- [SPARK-12941](#) - [SQL][MASTER] Spark-SQL JDBC Oracle dialect fails to map string datatypes to Oracle VARCHAR datatype
- [SPARK-12966](#) - [SQL] ArrayType(DecimalType) support in Postgres JDBC
- [SPARK-13112](#) - [CORE] Make sure RegisterExecutorResponse arrive before LaunchTask
- [SPARK-13242](#) - [SQL] codegen fallback in case-when if there many branches
- [SPARK-13328](#) - [CORE] Poor read performance for broadcast variables with dynamic resource allocation
- [SPARK-13566](#) - [CORE] Avoid deadlock between BlockManager and Executor Thread
- [SPARK-13958](#) - Executor OOM due to unbounded growth of pointer array

- [SPARK-14204](#) - [SQL] register driverClass rather than user-specified class
- [SPARK-14391](#) - [LAUNCHER] Fix launcher communication test, take 2.
- [SPARK-14963](#) - [MINOR][YARN] Fix typo in YarnShuffleService recovery file name. Using recoveryPath if NM recovery is enabled
- [SPARK-15165](#) - [SPARK-15205] [SQL] Introduce place holder for comments in generated code
- [SPARK-16044](#) - [SQL] Backport input_file_name() for data source based on NewHadoopRDD to branch 1.6
- [SPARK-16106](#) - [CORE] TaskSchedulerImpl should properly track executors added to existing hosts
- [SPARK-16230](#) - [CORE] CoarseGrainedExecutorBackend to self kill if there is an exception while creating an Executor
- [SPARK-16505](#) - [YARN] Optionally propagate error during shuffle service startup.
- [SPARK-16625](#) - [SQL] General data types to be mapped to Oracle
- [SPARK-16711](#) - YarnShuffleService doesn't re-init properly on YARN rolling upgrade
- [SPARK-16873](#) - [CORE] Fix SpillReader NPE when spillFile has no data
- [SPARK-17171](#) - [WEB UI] DAG will list all partitions in the graph
- [SPARK-17245](#) - [SQL][BRANCH-1.6] Do not rely on Hive's session state to retrieve HiveConf
- [SPARK-17433](#) - YarnShuffleService doesn't handle moving credentials levelDb
- [SPARK-17465](#) - [SPARK CORE] Inappropriate memory management in `org.apache.spark.storage.MemoryStore` may lead to memory leak
- [SPARK-17611](#) - [YARN][TEST] Make shuffle service test really test auth.
- [SPARK-17644](#) - [CORE] Do not add failedStages when abortStage for fetch failure
- [SPARK-17696](#) - [SPARK-12330][CORE] Partial backport of to branch-1.6.
- [SPARK-18750](#) - [YARN] Avoid using "mapValues" when allocating containers.
- [SPARK-19178](#) - [SQL][Backport-to-1.6] convert string of large numbers to int should return null
- [SPARK-19263](#) - DAGScheduler should avoid sending conflicting task set.
- [SPARK-19537](#) - Move pendingPartitions to ShuffleMapStage.
- [SQOOP-2349](#) - Add command line option for setting transaction isolation levels for metadata queries
- [SQOOP-2561](#) - Special Character removal from Column name as avro data results in duplicate column and fails the import
- [SQOOP-2846](#) - Sqoop Export with update-key failing for avro data file
- [SQOOP-2884](#) - Document --temporary-rootdir
- [SQOOP-2896](#) - Sqoop exec job fails with SQLException Access denied for user
- [SQOOP-2906](#) - Optimization of AvroUtil.toAvroIdentifier
- [SQOOP-2909](#) - Oracle related ImportTest fails after SQOOP-2737
- [SQOOP-2911](#) - Fix failing HCatalogExportTest caused by SQOOP-2863
- [SQOOP-2915](#) - Fixing Oracle related unit tests
- [SQOOP-2920](#) - sqoop performance deteriorates significantly on wide datasets; sqoop 100% on cpu
- [SQOOP-2950](#) - Sqoop trunk has consistent UT failures - need fixing
- [SQOOP-2952](#) - Fixing bug
- [SQOOP-2971](#) - OraOop does not close connections properly
- [SQOOP-2983](#) - OraOop export has degraded performance with wide tables
- [SQOOP-2986](#) - Add validation check for --hive-import and --incremental lastmodified
- [SQOOP-2990](#) - Sqoop(oracle) export [updateTableToOracle] with "--update-mode allowinsert" : app fails with java.sql.SQLException: Missing IN or OUT parameter at index
- [SQOOP-2995](#) - backward incompatibility introduced by Custom Tool options
- [SQOOP-2999](#) - Sqoop ClassNotFoundException (org.apache.commons.lang3.StringUtils) is thrown when executing Oracle direct import map task
- [SQOOP-3013](#) - Configuration "tmpjars" is not checked for empty strings before passing to MR
- [SQOOP-3021](#) - ClassWriter fails if a column name contains a backslash character
- [SQOOP-3028](#) - Include stack trace in the logging of exceptions in ExportTool
- [SQOOP-3034](#) - HBase import should fail fast if using anything other than as-textfile
- [SQOOP-3053](#) - Create a cmd line argument for sqoop.throwOnError and use it through SqoopOptions

- [SQOOP-3055](#) - Fixing MySQL tests failing due to ignored test inputs/configuration
- [SQOOP-3057](#) - Fixing 3rd party Oracle tests failing due to invalid case of column names
- [SQOOP-3066](#) - Introduce an option + env variable to enable/disable SQOOP-2737 feature
- [SQOOP-3068](#) - Enhance error (tool.ImportTool: Encountered IOException running import job: java.io.IOException: Expected schema) to suggest workaround
- [SQOOP-3069](#) - Get OracleExportTest#testUpsertTestExport in line with SQOOP-3066
- [SQOOP-3071](#) - Fix OracleManager to apply localTimeZone correctly in case of Date objects too
- [SQOOP-3072](#) - Reenable escaping in ImportTest#testProductWithWhiteSpaceImport for proper execution
- [SQOOP-3081](#) - use OracleEscapeUtils.escapeldentifier in OracleUpsertOutputFormat instead of inline appending quotes
- [SQOOP-3123](#) - Introduce escaping logic for column mapping parameters (same what Sqoop already uses for the DB column names), thus special column names (e.g. containing '#' character) and mappings related to those columns can be in the same format (thus not confusing the end users), and also eliminates the related AVRO format clashing issues.
- [SQOOP-3124](#) - Fix ordering in column list query of PostgreSQL connector to reflect the logical order instead of adhoc ordering
- [SQOOP-3159](#) - Sqoop (export + --table) with Oracle table_name having '\$' fails with error

Issues Fixed in CDH 5.8.4

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.8.4:

- [AVRO-1943](#) - Unreliable test: TestNettyServerWithCompression.testConnectionsCount
- [CRUNCH-592](#) - Job fails for null ByteBuffer value in Avro tables
- [FLUME-2171](#) - Add Interceptor to remove headers from event
- [FLUME-2812](#) - Fix semaphore leak causing java.lang.Error: Maximum permit count exceeded in MemoryChannel
- [FLUME-2889](#) - Fixes to DateTime computations
- [FLUME-2997](#) - Fix unreliable test in SpillableMemoryChannel
- [FLUME-2999](#) - Kafka channel and sink should enable statically assigned partition per event via header
- [FLUME-3002](#) - Fix tests in TestBucketWriter
- [FLUME-3003](#) - Fix unreliable testSourceCounter in TestSyslogUdpSource
- [FLUME-3020](#) - Improve HDFS Sink escape sequence substitution
- [FLUME-3027](#) - Change Kafka Channel to clear offsets map after commit
- [FLUME-3031](#) - Change sequence source to reset its counter for event body on channel exception
- [HADOOP-7930](#) - Kerberos relogin interval in UserGroupInformation should be configurable
- [HADOOP-10300](#) - Deferred sending of call responses allowed
- [HADOOP-11031](#) - Design Document for Credential Provider API
- [HADOOP-12453](#) - Support decoding KMS Delegation Token with its own Identifier
- [HADOOP-12483](#) - Maintain wrapped SASL ordering for postponed IPC responses
- [HADOOP-12537](#) - S3A to support Amazon STS temporary credentials
- [HADOOP-12655](#) - TestHttpServer.testBindAddress bind port range is wider than expected
- [HADOOP-12723](#) - S3A to add ability to plug in any AWSCredentialsProvider
- [HADOOP-12973](#) - Make DU pluggable
- [HADOOP-12974](#) - Create a CachingGetSpaceUsed implementation that uses df
- [HADOOP-12975](#) - Add jitter to CachingGetSpaceUsed's thread
- [HADOOP-13034](#) - Log message about input options in distcp lacks some items
- [HADOOP-13072](#) - WindowsGetSpaceUsed constructor should be public
- [HADOOP-13317](#) - Add logs to KMS server-side to improve supportability
- [HADOOP-13590](#) - Retry until TGT expires even if the UGI renewal thread encountered exception
- [HADOOP-13641](#) - Update UGI#spawnAutoRenewalThreadForUserCreds to reduce indentation
- [HADOOP-13669](#) - Addendum patch 2 for KMS Server should log exceptions before throwing

- [HADOOP-13669](#) - Addendum patch for KMS Server should log exceptions before throwing
- [HADOOP-13669](#) - KMS Server should log exceptions before throwing
- [HADOOP-13693](#) - Remove the message about HTTP OPTIONS in SPNEGO initialization message from kms audit log
- [HADOOP-13838](#) - KMSTokenRenewer should close providers
- [HDFS-4176](#) - EditLogTailer should call rollEdits with a timeout
- [HDFS-6962](#) - ACLs inheritance conflict with umaskmode
- [HDFS-7413](#) - Some unit tests should use NameNodeProtocols instead of FSNameSystem
- [HDFS-7964](#) - Add support for async edit logging
- [HDFS-8709](#) - Clarify automatic sync in FSEditLog#logEdit
- [HDFS-9038](#) - DFS reserved space is erroneously counted towards non-DFS used
- [HDFS-9630](#) - DistCp minor refactoring and clean up
- [HDFS-9638](#) - Improve DistCp Help and documentation
- [HDFS-9820](#) - Improve distcp to support efficient restore to an earlier snapshot
- [HDFS-10216](#) - Distcp -diff throws exception when handling relative path
- [HDFS-10298](#) - Document the usage of distcp -diff option
- [HDFS-10312](#) - Large block reports may fail to decode at NameNode due to 64 MB protobuf maximum length restriction
- [HDFS-10313](#) - Distcp need to enforce the order of snapshot names passed to -diff
- [HDFS-10336](#) - TestBalancer failing intermittently because of not resetting UserGroupInformation completely
- [HDFS-10397](#) - Distcp should ignore -delete option if -diff option is provided instead of exiting
- [HDFS-10556](#) - DistCpOptions should be validated automatically
- [HDFS-10609](#) - Uncaught InvalidEncryptionKeyException during pipeline recovery may abort downstream applications
- [HDFS-10652](#) - Add a unit test for HDFS-4660
- [HDFS-10722](#) - Fix race condition in TestEditLog#testBatchedSyncWithClosedLogs
- [HDFS-10760](#) - DataXceiver#run() should not log InvalidToken exception as an error
- [HDFS-10763](#) - Open files can leak permanently due to inconsistent lease update
- [HDFS-11012](#) - Unnecessary INFO logging on DFSClients for InvalidToken
- [HDFS-11040](#) - Add documentation for HDFS-9820 distcp improvement
- [HDFS-11056](#) - Concurrent append and read operations lead to checksum error
- [HDFS-11160](#) - VolumeScanner reports write-in-progress replicas as corrupt incorrectly
- [HDFS-11229](#) - HDFS-11056 failed to close meta file
- [HDFS-11275](#) - Check groupEntryIndex and throw a helpful exception on failures when removing ACL
- [MAPREDUCE-6571](#) - JobEndNotification info logs are missing in AM container syslog
- [MAPREDUCE-6633](#) - AM should retry map attempts if the reduce task encounters compression related errors
- [MAPREDUCE-6728](#) - Give fetchers hint when ShuffleHandler rejects a shuffling connection
- [MAPREDUCE-6763](#) - Shuffle server listen queue is too small
- [MAPREDUCE-6798](#) - Fix intermittent failure of TestJobHistoryParsing.testJobHistoryMethods
- [MAPREDUCE-6801](#) - Fix unreliable TestKill.testKillJob
- [MAPREDUCE-6817](#) - The format of job start time in JHS is different from submit and finish times
- [YARN-3601](#) - Fix UT TestRMFailover.testRMWebAppRedirect
- [YARN-3654](#) - ContainerLogsPage web UI should not have meta-refresh
- [YARN-3722](#) - Merge multiple TestWebAppUtils into o.a.h.yarn.webapp.util.TestWebAppUtils
- [YARN-4004](#) - Container-executor should print output of docker logs if the docker container exits with non-0 exit status
- [YARN-4017](#) - Container-executor overuses PATH_MAX
- [YARN-4092](#) - Fixed UI redirection to print useful messages when both RMs are in standby mode (Addendum)
- [YARN-4245](#) - Generalize config file handling in container-executor
- [YARN-4255](#) - Container-executor does not clean up docker operation command files
- [YARN-4544](#) - All the log messages about rolling monitoring interval are shown with WARN level

- [YARN-4556](#) - TestFifoScheduler.testResourceOverCommit fails
- [YARN-4820](#) - ResourceManager web redirects in HA mode drops query parameters
- [YARN-5001](#) - Aggregated Logs root directory is created with wrong group if nonexistent
- [YARN-5136](#) - Error in handling event type APP_ATTEMPT_REMOVED to the scheduler
- [YARN-5246](#) - NMWebAppFilter web redirects drop query parameters
- [YARN-5704](#) - Provide config knobs to control enabling/disabling new/work in progress features in container-executor
- [YARN-5837](#) - NPE when getting node status of a decommissioned node after an RM restart
- [YARN-5862](#) - TestDiskFailures.testLocalDirsFailures failed
- [YARN-5890](#) - FairScheduler should log information about AM-resource-usage and max-AM-share for queues
- [HBASE-15324](#) - Jitter may cause desiredMaxFileSize overflow in ConstantSizeRegionSplitPolicy and trigger unexpected split
- [HBASE-15430](#) - Failed taking snapshot - Manifest proto-message too large
- [HBASE-16146](#) - Counter performance is expensive
- [HBASE-16172](#) - Unify the retry logic in ScannerCallableWithReplicas and RpcRetryingCallerWithReadReplicas
- [HBASE-16270](#) - Handle duplicate clearing of snapshot in region replicas
- [HBASE-16345](#) - RpcRetryingCallerWithReadReplicas#call() should catch some RegionServer Exceptions
- [HBASE-16824](#) - Writer.flush() can be called on already closed streams in WAL roll
- [HBASE-16841](#) - Data loss in MOB files after cloning a snapshot and deleting that snapshot
- [HBASE-17058](#) - Lower epsilon used for jitter verification from HBASE-15324
- [HBASE-17072](#) - CPU usage starts to climb up to 90-100% when using G1GC
- [HBASE-17241](#) - Avoid compacting already compacted mob files with _del files
- [HBASE-17452](#) - Failed taking snapshot - region Manifest proto-message too large
- [HIVE-10384](#) - BackportRetryingMetaStoreClient does not retry wrapped TTransportExceptions
- [HIVE-11849](#) - NPE in HiveHBaseTableShapshotInputFormat in query with just count(*)
- [HIVE-12077](#) - MSCK Repair table should fix partitions in batches
- [HIVE-12465](#) - Hive might produce wrong results when (outer) joins are merged
- [HIVE-12619](#) - Switching the field order within an array of structs causes the query to fail
- [HIVE-12780](#) - Fix the output of the history command in Beeline
- [HIVE-12789](#) - Fix output twice in the history command of Beeline
- [HIVE-12891](#) - Hive fails when java.io.tmpdir is set to a relative location
- [HIVE-12976](#) - MetaStoreDirectSql doesn't batch IN lists in all cases
- [HIVE-13129](#) - CliService leaks HMS connection
- [HIVE-13149](#) - Remove some unnecessary HMS connections from HS2
- [HIVE-13149](#) - Remove some unnecessary HMS connections from HS2
- [HIVE-13381](#) - Timestamp and date should have precedence in type hierarchy than string group
- [HIVE-13429](#) - Tool to remove dangling scratch dir
- [HIVE-13539](#) - HiveHFileOutputFormat searching the wrong directory for HFiles
- [HIVE-13866](#) - Flatten callstack for directSQL errors
- [HIVE-13895](#) - HoS start-up overhead in yarn-client mode
- [HIVE-13997](#) - Insert overwrite directory doesn't overwrite existing files
- [HIVE-14137](#) - Hive on Spark throws FileAlreadyExistsException for jobs with multiple empty tables
- [HIVE-14173](#) - NPE was thrown after enabling direct.sql in the middle of session
- [HIVE-14421](#) - FS.deleteOnExit holds references to _tmp_space.db files
- [HIVE-14762](#) - Add logging while removing scratch space
- [HIVE-14799](#) - Query operation are not thread safe during its cancellation
- [HIVE-14817](#) - Shutdown the SessionManager timeoutChecker thread properly upon shutdown
- [HIVE-15054](#) - Hive insertion query execution fails on Hive on Spark
- [HIVE-15061](#) - Metastore types are sometimes case sensitive
- [HIVE-15090](#) - Temporary DB failure can stop ExpiredTokenRemover thread
- [HIVE-15231](#) - Query on view with CTE and alias fails with "table not found error"

- [HIVE-15291](#) - Comparison of timestamp fails if only date part is provided
- [HIVE-15410](#) - WebHCat supports get/set table property with its name containing period and hyphen
- [HIVE-15551](#) - Memory leak in directsql for MySQL with BoneCP specific initialization
- [HUE-4466](#) - [security] deliver csrftoken cookie with secure bit set if possible.
- [HUE-4546](#) - Auto-strip invalid characters from name field of converted docs
- [HUE-4747](#) - [editor] Download form should be submitted to a new tab otherwise the snippet gets closed
- [HUE-5028](#) - [oozie] User can't edit shared WF with modify permissions in new editor mode
- [HUE-5050](#) - [core] Logout fails for local login when multiple backends are used
- [HUE-5154](#) - [oozie] Create an new oozie workflow throws server error 500
- [HUE-5161](#) - [security] Speed up roles rendering
- [HUE-5163](#) - [security] Speed up initial page rendering
- [HUE-5166](#) - [impala] Handle empty session properties during upgrades
- [HUE-5218](#) - [search] Validate dashboard sharing works
- [HUE-5295](#) - [desktop] Avoid microsecond comparison for last_modified field MySQL < 5.6 doesn't support microsecond precision. <https://code.djangoproject.com/ticket/19716>
- [HUE-5295](#) - [desktop] Do not change the last_modified field when migrating history queries
- [HUE-5305](#) - [home] Fix empty share document modal and improve sharing UX
- [HUE-5310](#) - [search] Use Doc2 modal in search_controller
- [HUE-5476](#) - [core] Fix TTL is_idle middleware check
- [HUE-5482](#) - [home] Handle multiple home/trash directories by merging them into one.
- [IMPALA-1702](#) - "invalidate metadata" can cause duplicate TableIds
- [IMPALA-3167](#) - Fix assignment of WHERE clause predicate through grouping aggregate and outer join
- [IMPALA-3314](#) - Fix Avro schema loading for partitioned tables
- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3575](#) - Add retry to backend connection request and rpc timeout
- [IMPALA-3682](#) - Don't retry unrecoverable socket creation errors
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3884](#) - Support TYPE_TIMESTAMP for HashTableCtx::CodegenAssignNullValue()
- [IMPALA-3949](#) - Log the error message in FileSystemUtil.copyToLocal()
- [IMPALA-3964](#) - Fix crash when a count(*) is performed on a nested collection.
- [IMPALA-3983](#) - Delete function jar resources after load
- [IMPALA-4037](#) - ChildQuery::Cancel() appears to violate lock ordering
- [IMPALA-4038](#) - Fix locking during query cancellation
- [IMPALA-4076](#) - Fix runtime filter sort compare method
- [IMPALA-4099](#) - Fix the error message while loading UDFs with no JARs
- [IMPALA-4120](#) - Incorrect results with LEAD() analytic function
- [IMPALA-4153](#) - Fix count(*) on all blank("") columns - test
- [IMPALA-4223](#) - Handle truncated file read from HDFS cache
- [IMPALA-4246](#) - SleepForMs() utility function has undefined behavior for > 1s
- [IMPALA-4336](#) - Cast expressions after unnesting union operands
- [IMPALA-4363](#) - Add Parquet timestamp validation
- [IMPALA-4391](#) - Fix dropped statuses in scanners
- [IMPALA-4423](#) - Correct but conservative implementation of Subquery.equals()
- [IMPALA-4433](#) - Always generate test data using the same time zone setting
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog. Fixes an issue where multiple long-running operations on the same catalog object (for example, a table) can block other catalog operations from making progress
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots
- [IMPALA-4391](#) - Fix dropped status in scanners
- [LUCENE-5889](#) - AnalyzingInfixSuggester should expose commit()
- [LUCENE-7564](#) - AnalyzingInfixSuggester should close its IndexWriter by default at the end of build()

- [PIG-3818](#) - PIG-2499 is accidentally reverted
- [PIG-5025](#) - Fix unreliable test failures in TestLoad.java
- [SENTRY-1265](#) - Sentry service should not require a TGT as it is not talking to other Kerberos services as a client
- [SENTRY-1313](#) - Database prefix is not honoured when executing grant statement
- [SPARK-12241](#) - [YARN] Improve failure reporting in Yarn client obtainTokenForHBase()
- [SPARK-12523](#) - [YARN] Support long-running of the Spark On HBase and hive meta store.
- [SPARK-12966](#) - [SQL] ArrayType(DecimalType) support in Postgres JDBC
- [SPARK-13566](#) - [CORE] Avoid deadlock between BlockManager and Executor Thread
- [SPARK-13958](#) - Executor OOM due to unbounded growth of pointer array in...
- [SPARK-14204](#) - [SQL] register driverClass rather than user-specified class
- [SPARK-16044](#) - [SQL] Backport input_file_name() for data source based on NewHadoopRDD to branch 1.6
- [SPARK-17245](#) - [SQL][BRANCH-1.6] Do not rely on Hive's session state to retrieve HiveConf
- [SPARK-17465](#) - [SPARK CORE] Inappropriate memory management in `org.apache.spark.storage.MemoryStore` may lead to memory leak
- [SQOOP-2349](#) - Add command line option for setting transaction isolation levels for metadata queries
- [SQOOP-2884](#) - Document argument overriding --temporary directory
- [SQOOP-2909](#) - Oracle related ImportTest fails after SQOOP-2737
- [SQOOP-2911](#) - Fix failing HCatalogExportTest caused by SQOOP-2863
- [SQOOP-2915](#) - Fixing Oracle related unit tests
- [SQOOP-2950](#) - Fix Sqoop trunk consistent UT failures
- [SQOOP-2952](#) - Row key not added into column family using --hbase-bulkload
- [SQOOP-2983](#) - OraOOP export has degraded performance with wide tables
- [SQOOP-2986](#) - Add validation check for --hive-import and --incremental lastmodified
- [SQOOP-2990](#) - Sqoop(oracle) export [updateTableToOracle] with "--update-mode allowinsert" : app fails with java.sql.SQLException: Missing IN or OUT parameter at index
- [SQOOP-3013](#) - Configuration "tmpjars" is not checked for empty strings before passing to MR
- [SQOOP-3028](#) - Include stack trace in the logging of exceptions in ExportTool
- [SQOOP-3034](#) - HBase import should fail fast if using anything other than as-textfile
- [SQOOP-3053](#) - Create a cmd line argument for sqoop.throwOnError and use it through SqoopOptions
- [SQOOP-3055](#) - Fixing MySQL tests failing due to ignored test inputs/configuration
- [SQOOP-3057](#) - Fixing 3rd party Oracle tests failing due to invalid case of column names
- [SQOOP-3066](#) - Introduce an option + env variable to enable/disable SQOOP-2737 feature
- [SQOOP-3069](#) - Get OracleExportTest#testUpsertTestExport in line with SQOOP-3066
- [SQOOP-3071](#) - Fix OracleManager to apply localTimeZone correctly in case of Date objects too
- [SQOOP-3072](#) - Re-enable escaping in ImportTest#testProductWithWhiteSpaceImport for proper execution
- [SQOOP-3081](#) - Use OracleEscapeUtils.escapeIdentifier in OracleUpsertOutputFormat instead of inline appending quotes
- [SQOOP-3124](#) - Fix ordering in column list query of PostgreSQL connector to reflect logical order (rather than ad hoc ordering)

Issues Fixed in CDH 5.8.3

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.8.3:

- [FLUME-2797](#) - Use SourceCounter for SyslogTcpSource
- [FLUME-2844](#) - SpillableMemoryChannel must start ChannelCounter
- [HADOOP-12548](#) - Read s3a credentials from a Credential Provider
- [HADOOP-13353](#) - LdapGroupsMapping getPassword should not return null when IOException throws
- [HADOOP-13526](#) - Add detailed logging in KMS for the authentication failure of proxy user
- [HADOOP-13558](#) - UserGroupInformation created from a Subject incorrectly tries to renew the Kerberos ticket
- [HADOOP-13579](#) - Fix source-level compatibility after HADOOP-11252

- [HADOOP-13638](#) - KMS should set UGI's Configuration object properly
- [HDFS-7415](#) - Move FSNameSystem.resolvePath() to FSDirectory
- [HDFS-7420](#) - Delegate permission checks to FSDirectory
- [HDFS-7463](#) - Simplify FSNamesystem#getBlockLocationsUpdateTimes
- [HDFS-7478](#) - Move org.apache.hadoop.hdfs.server.namenode.NNConf to FSNamesystem
- [HDFS-7517](#) - Remove redundant non-null checks in FSNamesystem#getBlockLocations
- [HDFS-8224](#) - Schedule a block for scanning if its metadata file is corrupt
- [HDFS-8269](#) - getBlockLocations() does not resolve the .reserved path and generates incorrect edit logs when updating the atime
- [HDFS-9601](#) - NNThroughputBenchmark.BlockReportStats should handle NotReplicatedYetException on adding block.
- [HDFS-9781](#) - FsDatasetImpl#getBlockReports can occasionally throw NullPointerException
- [HDFS-10641](#) - TestBlockManager#testBlockReportQueueing fails intermittently
- [HDFS-10879](#) - TestEncryptionZonesWithKMS#testReadWrite fails intermittently
- [HDFS-10962](#) - TestRequestHedgingProxyProvider fails intermittently
- [HDFS-10963](#) - Reduce log level when network topology cannot find enough datanodes
- [MAPREDUCE-6628](#) - Potential memory leak in CryptoOutputStream
- [MAPREDUCE-6641](#) - TestTaskAttempt fails in trunk
- [MAPREDUCE-6718](#) - Add progress log to JHS during startup
- [MAPREDUCE-6771](#) - RMContainerAllocator sends container diagnostics event after corresponding completion event
- [YARN-4940](#) - yarn node -list -all fails if RM starts with decommissioned node
- [HBASE-15856](#) - Do not cache unresolved addresses for connections
- [HBASE-16294](#) - hbck reporting "No HDFS region dir found" for replicas
- [HBASE-16699](#) - Overflows in AverageIntervalRateLimiter's refill() and getWaitInterval()
- [HBASE-16767](#) - Mob compaction needs to clean up files in /hbase/mobdir/.tmp and /hbase/mobdir/.tmp/.bulkload when running into IO exceptions
- [HIVE-10965](#) - Direct SQL for stats fails in 0-column case
- [HIVE-12083](#) - HIVE-10965 introduces thrift error if partNames or colNames are empty
- [HIVE-12475](#) - Parquet schema evolution within array<struct<>> does not work
- [HIVE-12785](#) - View with union type and UDF to the struct is broken
- [HIVE-13058](#) - Add session and operation_log directory deletion messages
- [HIVE-13198](#) - Authorization issues with cascading views
- [HIVE-13237](#) - Select parquet struct field with upper case throws NPE
- [HIVE-13620](#) - Merge llap branch work to master
- [HIVE-13625](#) - Hive Prepared Statement when executed with escape characters in parameter fails
- [HIVE-13645](#) - Beeline needs null-guard around hiveVars and hiveConfVars read
- [HIVE-14296](#) - Session count is not decremented when HS2 clients do not shutdown cleanly
- [HIVE-14383](#) - SparkClientImpl should pass principal and keytab to spark-submit instead of calling kinit explicitly
- [HIVE-14715](#) - Hive throws NumberFormatException with query with Null value
- [HIVE-14743](#) - ArrayIndexOutOfBoundsException - HBASE-backed views' query with JOINS
- [HIVE-14784](#) - Operation logs are disabled automatically if the parent directory does not exist.
- [HIVE-14805](#) - Subquery inside a view will have the object in the subquery as the direct input
- [HUE-4064](#) - Format creation and update date on the table details popover
- [HUE-4138](#) - Last modified time of a saved query is not in the correct timezone
- [HUE-4141](#) - Graph breaks for external workflows when there is more than one kill node
- [HUE-4804](#) - Download function of HTML widget breaks the display
- [HUE-4809](#) - Add trustore parameters only if SSL is turned on
- [HUE-4809](#) - Only add trustore paths when they are actually existing
- [HUE-4810](#) - Fix tests by setting data to valid JSON type

- [HUE-4871](#) - An unprivileged user can enumerate users
- [HUE-4891](#) - An unprivileged user can list document items
- [HUE-4916](#) - Truncate last name to 30 chars on ldap import
- [HUE-4968](#) - Remove access to /oozie/import_workflow when v2 is enabled
- [HUE-4994](#) - Consider default path for decision nodes in dashboard graph
- [HUE-5041](#) - Hue export large file to HDFS does not work on non-default database
- [IMPALA-1619](#) - Support 64-bit allocations
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3751](#) - Fix clang build errors and warnings
- [IMPALA-4135](#) - Thrift threaded server times-out connections during high load
- [IMPALA-4170](#) - Fix identifier quoting in COMPUTE INCREMENTAL STATS
- [IMPALA-4180](#) - Synchronize accesses to RuntimeState::reader_contexts_
- [IMPALA-4196](#) - Cross compile bit-byte-functions
- [IMPALA-4237](#) - Fix materialization of 4-byte decimals in data source scan node
- [OOZIE-1814](#) - Oozie should mask any passwords in logs and REST interfaces
- [SQLR-9310](#) - PeerSync fails on a node restart due to IndexFingerPrint mismatch
- [SPARK-12009](#) - Avoid reallocating YARN container when driver wants to stop all Executors
- [SPARK-12392](#) - Optimize a location order of broadcast blocks by considering preferred local hosts
- [SPARK-12941](#) - Spark-SQL JDBC Oracle dialect fails to map string datatypes to Oracle VARCHAR datatype mapping
- [SPARK-12941](#) - Spark-SQL JDBC Oracle dialect fails to map string datatypes to Oracle VARCHAR datatype
- [SPARK-13328](#) - Poor read performance for broadcast variables with dynamic resource allocation
- [SPARK-16625](#) - General data types to be mapped to Oracle
- [SPARK-16711](#) - YarnShuffleService doesn't re-init properly on YARN rolling upgrade
- [SPARK-17171](#) - DAG will list all partitions in the graph
- [SPARK-17433](#) - YarnShuffleService doesn't handle moving credentials levelDb
- [SPARK-17611](#) - Make shuffle service test really test authentication
- [SPARK-17644](#) - Do not add failedStages when abortStage for fetch failure
- [SPARK-17696](#) - Partial backport of to branch-1.6.
- [SQOOP-3021](#) - ClassWriter fails if a column name contains a backslash character

Issues Fixed in CDH 5.8.2

CDH 5.8.2 fixes the following issues.

Kerberized HS2 with LDAP authentication fails in a multi-domain LDAP case

In CDH 5.7, Hive introduced a feature to support HS2 with Kerberos plus LDAP authentication; but it broke compatibility with multi-domain LDAP cases on CDH 5.7.x and C5.8.x versions.

Affected Versions: CDH 5.7.1, CDH 5.8.0, and CDH 5.8.1

Fixed in Versions: CDH 5.7.2 and higher, CDH 5.8.2 and higher

Bug: [HIVE-13590](#).

Workaround: None.

Apache Oozie

Oozie Web Console returns 500 error when Oozie server runs on JDK 8u75 or higher

Bug: [OOZIE-2533](#)

The Oozie Web Console returns a 500 error when the Oozie server is running on JDK 8u75 and higher. The Oozie server still functions, and you can use the Oozie command line, REST API, Java API, or the Hue Oozie Dashboard to review status of those jobs.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.8.2:

- [FLUME-1899](#) - Make SpoolDir work with subdirectories
- [FLUME-2652](#) - Documented transaction handling semantics incorrect in developer guide.
- [FLUME-2901](#) - Document Kerberos setup for Kafka channel
- [FLUME-2910](#) - AsyncHBaseSink: Failure callbacks should log the exception that caused them
- [FLUME-2913](#) - Don't strip SLF4J from imported classpaths
- [FLUME-2918](#) - Speed up TaildirSource on directories with many files
- [FLUME-2922](#) - Sync SequenceFile.Writer before calling hflush
- [FLUME-2923](#) - Bump asynchbase version to 1.7.0
- [FLUME-2934](#) - Document new cachePatternMatching option for TaildirSource
- [FLUME-2935](#) - Bump java target version to 1.7
- [FLUME-2948](#) - docs: Fix parameters on Replicating Channel Selector example
- [FLUME-2954](#) - Make raw data appearing in log messages explicit
- [FLUME-2963](#) - FlumeUserGuide: Fix error in Kafka Source properties table
- [FLUME-2972](#) - Handle offset migration in the new Kafka Channel
- [FLUME-2975](#) - docs: Fix NetcatSource example
- [FLUME-2982](#) - Add localhost escape sequence to HDFS sink
- [FLUME-2983](#) - Handle offset migration in the new Kafka Source
- [HADOOP-8436](#) - NPE In getLocalPathForWrite (path, conf) when the required context item is not configured
- [HADOOP-8437](#) - getLocalPathForWrite should throw IOException for invalid paths
- [HADOOP-8934](#) - Shell command ls should include sort options (Jonathan Allen via aw)
- [HADOOP-8934](#) - Shell command ls should include sort options
- [HADOOP-10048](#) - LocalDirAllocator should avoid holding locks while accessing the filesystem
- [HADOOP-10971](#) - Add -C flag to make `hadoop fs -ls` print filenames only
- [HDFS-10512](#) - VolumeScanner can terminate due to NPE in DataNode.reportBadBlocks.
- [HADOOP-11361](#) - Fix a race condition in MetricsSourceAdapter.updateJmxCache.
- [HADOOP-11469](#) - KMS should skip default.key.acl and whitelist.key.acl when loading key acl.
- [HADOOP-11901](#) - BytesWritable fails to support 2G chunks due to integer overflow
- [HADOOP-12252](#) - LocalDirAllocator should not throw NPE with empty string configuration
- [HADOOP-12609](#) - Fix intermittent failure of TestDecayRpcScheduler.
- [HADOOP-12659](#) - Incorrect usage of config parameters in token manager of KMS
- [HADOOP-12963](#) - Allow using path style addressing for accessing the s3 endpoint.
- [HADOOP-13079](#) - Add -q option to ls to print ? instead of non-printable characters
- [HADOOP-13132](#) - Handle ClassCastException on AuthenticationException in LoadBalancingKMSSClientProvider
- [HADOOP-13155](#) - Implement TokenRenewer to renew and cancel delegation tokens in KMS
- [HADOOP-13251](#) - Authenticate with Kerberos credentials when renewing KMS delegation token
- [HADOOP-13255](#) - KMSSClientProvider should check and renew tgt when doing delegation token operations.
- [HADOOP-13263](#) - Reload cached groups in background after expiry.
- [HADOOP-13270](#) - BZip2CompressionInputStream finds the same compression marker twice in corner case, causing duplicate data blocks
- [HADOOP-13381](#) - KMS clients should use KMS Delegation Tokens from current UGI
- [HADOOP-13437](#) - KMS should reload whitelist and default key ACLs when hot-reloading [HADOOP-13457](#) - Remove hardcoded absolute path for shell executable.
- [HADOOP-13487](#) - Hadoop KMS should load old delegation tokens from Zookeeper on startup
- [HDFS-4210](#) - Throw helpful exception when DNS entry for JournalNode cannot be resolved
- [HDFS-6434](#) - Default permission for creating file should be 644 for WebHdfs/HttpFS
- [HDFS-7597](#) - DelegationTokenIdentifier should cache the TokenIdentifier to UGI mapping
- [HDFS-8581](#) - ContentSummary on / skips further counts on yielding lock
- [HDFS-8829](#) - Make SO_RCVBUF and SO_SNDBUF size configurable for DataTransferProtocol sockets and allow configuring auto-tuning
- [HDFS-8897](#) - Balancer should handle fs.defaultFS trailing slash in HA

- [HDFS-9085](#) - Show renewer information in DelegationTokenIdentifier#toString
- [HDFS-9137](#) - DeadLock between DataNode#refreshVolumes and BPOfferService#registrationSucceeded.
- [HDFS-9141](#) - Thread leak in Datanode#refreshVolumes.
- [HDFS-9259](#) - Make SO_SNDBUF size configurable at DFSClient side for hdfs write scenario.
- [HDFS-9276](#) - Failed to Update HDFS Delegation Token for long running application in HA mode
- [HDFS-9365](#) - Balancer does not work with the HDFS-6376 HA setup.
- [HDFS-9461](#) - DiskBalancer: Add Report Command
- [HDFS-9466](#) - TestShortCircuitCache#testDataXceiverCleansUpSlotsOnFailure is unreliable
- [HDFS-9700](#) - DFSClient and DFSOutputStream should set TCP_NODELAY on sockets for DataTransferProtocol
- [HDFS-9732](#) - Improve DelegationTokenIdentifier.toString() for better logging
- [HDFS-9805](#) - Add server-side configuration for enabling TCP_NODELAY for DataTransferProtocol and default it to true
- [HDFS-9906](#) - Remove spammy log spew when a datanode is restarted.
- [HDFS-9939](#) - Increase DecompressorStream skip buffer size
- [HDFS-9958](#) - BlockManager#createLocatedBlocks can throw NPE for corruptBlocks on failed storages
- [HDFS-10270](#) - TestJMXGet:testNameNode() fails
- [HDFS-10381](#) - , DataStreamer DataNode exclusion log message should be warning.
- [HDFS-10403](#) - DiskBalancer: Add cancel command
- [HDFS-10457](#) - DataNode should not auto-format block pool directory if VERSION is missing.
- [HDFS-10481](#) - HTTPFS server should correctly impersonate as end user to open file
- [HDFS-10500](#) - Diskbalancer: Print out information when a plan is not generated
- [HDFS-10501](#) - DiskBalancer: Use the default datanode port if port is not provided
- [HDFS-10516](#) - Fix bug when warming up EDEK cache of more than one encryption zone
- [HDFS-10517](#) - DiskBalancer: Support help command
- [HDFS-10525](#) - Fix NPE in CacheReplicationMonitor#rescanCachedBlockMap
- [HDFS-10541](#) - Diskbalancer: When no actions in plan, error message says "Plan was generated more than 24 hours ago"
- [HDFS-10544](#) - Balancer doesn't work with IPFailoverProxyProvider.
- [HDFS-10552](#) - DiskBalancer "-query" results in NPE if no plan for the node
- [HDFS-10559](#) - DiskBalancer: Use SHA1 for Plan ID
- [HDFS-10567](#) - Improve plan command help message
- [HDFS-10588](#) - False alarm in datanode log - ERROR - Disk Balancer is not enabled
- [HDFS-10598](#) - DiskBalancer does not execute multi-steps plan
- [HDFS-10600](#) - PlanCommand#getThresholdPercentage should not use throughput value.
- [HDFS-10643](#) - Namenode should use loginUser(hdfs) to generateEncryptedKey
- [HDFS-10681](#) - DiskBalancer: query command should report Plan file path apart from PlanID.
- [HDFS-10822](#) - Log DataNodes in the write pipeline. John Zhuge via Lei Xu
- [MAPREDUCE-4784](#) - TestRecovery occasionally fails
- [MAPREDUCE-6359](#) - In RM HA setup, Cluster tab links populated with AM hostname instead of RM
- [MAPREDUCE-6442](#) - Stack trace is missing when error occurs in client protocol provider's constructor Contributed by Chang Li.
- [MAPREDUCE-6473](#) - Revert "Revert "Job submission can take a long time during Cluster initialization
- [MAPREDUCE-6473](#) - Revert "Job submission can take a long time during Cluster initialization
- [MAPREDUCE-6473](#) - Job submission can take a long time during Cluster initialization
- [MAPREDUCE-6670](#) - TestJobListCache#testEviction sometimes fails on Windows with timeout
- [MAPREDUCE-6680](#) - JHS UserLogDir scan algorithm sometime could skip directory with update in CloudFS (Azure FileSystem, S3, etc
- [MAPREDUCE-6738](#) - TestJobListCache.testAddExisting failed intermittently in slow VM testbed
- [MAPREDUCE-6761](#) - Regression when handling providers - invalid configuration ServiceConfiguration causes Cluster initialization failure

- [YARN-2605](#) - [RM HA] Rest api endpoints doing redirect incorrectly.
- [YARN-2977](#) - Fixed intermittent TestNMClient failure.
- [YARN-4411](#) - RMAAppAttemptImpl#createApplicationAttemptReport throws IllegalArgumentException
- [YARN-4459](#) - container-executor should only kill process groups
- [YARN-4866](#) - FairScheduler: AMs can consume all vcores leading to a livelock when using FAIR policy.
- [YARN-4878](#) - Expose scheduling policy and max running apps over JMX for Yarn queues.
- [YARN-4989](#) - TestWorkPreservingRMRestart#testCapacitySchedulerRecovery fails intermittently
- [YARN-5048](#) - DelegationTokenRenewer#skipTokenRenewal may throw NPE
- [YARN-5077](#) - Fix FSLeafQueue#getFairShare() for queues with zero fairshare.
- [YARN-5107](#) - TestContainerMetrics fails.
- [YARN-5272](#) - Handle queue names consistently in FairScheduler.
- [YARN-5608](#) - TestAMRMClient.setup() fails with ArrayOutOfBoundsException
- [HBASE-14644](#) - Region in transition metric is broken -- addendum
- [HBASE-14644](#) - Region in transition metric is broken
- [HBASE-14818](#) - user_permission does not list namespace permissions
- [HBASE-14963](#) - Remove use of Guava Stopwatch from HBase client code
- [HBASE-15465](#) - userPermission returned by getUserPermission() for the selected namespace does not have namespace set
- [HBASE-15496](#) - Throw RowTooBigException only for user scan/get
- [HBASE-15621](#) - Suppress Hbase SnapshotHFile cleaner error messages when a snapshot is going on
- [HBASE-15683](#) - Min latency in latency histograms are emitted as Long.MAX_VALUE
- [HBASE-15698](#) - Increment TimeRange not serialized to server
- [HBASE-15746](#) - Remove extra RegionCoprocessor preClose() in RSRpcServices#closeRegion
- [HBASE-15808](#) - Reduce potential bulk load intermediate space usage and waste
- [HBASE-15872](#) - Split TestWALProcedureStore
- [HBASE-15873](#) - ACL for snapshot restore / clone is not enforced
- [HBASE-15925](#) - provide default values for hadoop compat module related properties that match default hadoop profile.
- [HBASE-16034](#) - Fix ProcedureTestingUtility#LoadCounter.setMaxProcid()
- [HBASE-16056](#) - Procedure v2 - fix master crash for FileNotFound
- [HBASE-16093](#) - Fix splits failed before creating daughter regions leave meta inconsistent
- [HBASE-16135](#) - PeerClusterZnode under rs of removed peer may never be deleted
- [HBASE-16194](#) - Should count in MSLAB chunk allocation into heap size change when adding duplicate cells
- [HBASE-16195](#) - Should not add chunk into chunkQueue if not using chunk pool in HeapMemStoreLAB
- [HBASE-16207](#) - can't restore snapshot without "Admin" permission
- [HBASE-16227](#) - [Shell] Column value formatter not working in scans. Tested : manually using shell.
- [HBASE-16284](#) - Unauthorized client can shutdown the cluster
- [HBASE-16288](#) - HFile intermediate block level indexes might recurse forever creating multi TB files.
- [HBASE-16319](#) - Fix TestCacheOnWrite after HBASE-16288.
- [HBASE-16317](#) - revert all ESAPI changes
- [HBASE-16318](#) - fail build while rendering velocity template if dependency license isn't in whitelist.
- [HBASE-16318](#) - consistently use the correct name for 'Apache License, Version 2.0'
- [HBASE-16321](#) - ensure no findbugs-jsr305
- [HBASE-16340](#) - exclude Xerces implementation jars from coming in transitively.
- [HBASE-16360](#) - TableMapReduceUtil addHBaseDependencyJars has the wrong class name for PrefixTreeCodec
- [HIVE-7443](#) - Fix HiveConnection to communicate with Kerberized Hive JDBC server and alternative JDKs
- [HIVE-10007](#) - Support qualified table name in analyze table compute statistics for columns
- [HIVE-10728](#) - deprecate unix_timestamp(void) and make it deterministic
- [HIVE-11243](#) - Changing log level in Utilities.getBaseWork
- [HIVE-11432](#) - Hive macro give same result for different arguments

- [HIVE-11487](#) - Add getNumPartitionsByFilter api in metastore api
- [HIVE-11747](#) - Unnecessary error log is shown when executing a "INSERT OVERWRITE LOCAL DIRECTORY" cmd in the embedded mode
- [HIVE-11827](#) - STORED AS AVRO fails SELECT COUNT(*) when empty
- [HIVE-11901](#) - StorageBasedAuthorizationProvider requires write permission on table for SELECT statements
- [HIVE-11980](#) - Follow up on HIVE-11696, exception is thrown from CTAS from the table with table-level serde is Parquet while partition-level serde is JSON
- [HIVE-12277](#) - Hive macro results on macro_duplicate.q different after adding ORDER BY
- [HIVE-12556](#) - Ctrl-C in beeline doesn't kill Tez query on HS2
- [HIVE-12635](#) - Hive should return the latest hbase cell timestamp as the row timestamp value
- [HIVE-13043](#) - Reload function has no impact to function registry
- [HIVE-13090](#) - Hive metastore crashes on NPE with ZooKeeperTokenStore
- [HIVE-13372](#) - Hive Macro overwritten when multiple macros are used in one column
- [HIVE-13462](#) - HiveResultSetMetaData.getPrecision() fails for NULL columns
- [HIVE-13590](#) - Kerberized HS2 with LDAP auth enabled fails in multi-domain LDAP case
- [HIVE-13704](#) - Don't call DistCp.execute() instead of DistCp.run()
- [HIVE-13736](#) - View's input/output formats are TEXT by default.
- [HIVE-13749](#) - Memory leak in Hive Metastore
- [HIVE-13884](#) - Disallow queries in HMS fetching more than a configured number of partitions
- [HIVE-13932](#) - Hive SMB Map Join with small set of LIMIT failed with NPE
- [HIVE-13953](#) - Issues in HiveLockObject equals method
- [HIVE-13991](#) - Union All on view fail with no valid permission on underneath table
- [HIVE-14006](#) - Hive query with UNION ALL fails with ArrayIndexOutOfBoundsException.
- [HIVE-14015](#) - SMB MapJoin failed for Hive on Spark when kerberized
- [HIVE-14055](#) - directSql - getting the number of partitions is broken
- [HIVE-14098](#) - Logging task properties, and environment variables might contain passwords
- [HIVE-14118](#) - Make the alter partition exception more meaningful
- [HIVE-14187](#) - JDOPersistenceManager objects remain cached if MetaStoreClient#close is not called
- [HIVE-14209](#) - Add some logging info for session and operation management
- [HIVE-14436](#) - Hive 1.2.1/Hitting "ql.Driver: FAILED: IllegalArgumentException Error: , expected at the end of 'decimal(9'" after enabling hive.optimize.skewjoin and with MR engine
- [HIVE-14457](#) - Partitions in encryption zone are still trashed though an exception is returned
- [HIVE-14519](#) - Multi insert query bug
- [HIVE-14538](#) - beeline throws exceptions with parsing hive config when using !sh statement
- [HIVE-14697](#) - Can not access kerberized HS2 Web UI
- [HUE-2689](#) - Sub-workflow submitted from coordinator gets parent workflow graph
- [HUE-2971](#) - Some links of a Fork can point to deleted nodes
- [HUE-3842](#) - HTTP 500 while emptying Hue 3.9 trash directory
- [HUE-3908](#) - [useradmin] Ignore (objectclass=*) filter when searching for LDAP users
- [HUE-3988](#) - Support schemaless collections
- [HUE-3999](#) - list_oozie_workflow page shouldn't break incase of bad json from oozie
- [HUE-4005](#) - Remove oozie.coord.application.path from properties when rerunning workflow
- [HUE-4006](#) - Create new deployment directory when coordinator or bundle is copied
- [HUE-4007](#) - Fix deployment_dir for the bundle in oozie example fixtures
- [HUE-4019](#) - Always fetch the logs on check status
- [HUE-4019](#) - Do not blank error on query with good syntax but invalid query
- [HUE-4021](#) - [libsolr] Allow customization of the Solr path in ZooKeeper
- [HUE-4023](#) - [useradmin] update AuthenticationForm to allow activated users to login
- [HUE-4078](#) - Drag & Drop hive queries shows queries from the trash
- [HUE-4087](#) - Unable to kill jobs with Resource Manager HA enabled

- [HUE-4092](#) - Can't type any / in the HDFS ACLs path input
- [HUE-4119](#) - Change list jobs call to POST
- [HUE-4129](#) - Long running query getting terminated when leaving the editor
- [HUE-4134](#) - [liboozie] Avoid logging truststore credentials
- [HUE-4145](#) - Older queries after upgrade do not provide direct save
- [HUE-4146](#) - Older saved queries defaults to default' DB
- [HUE-4148](#) - Improve import testing of beeswax queries to notebook format
- [HUE-4153](#) - Report last seen progress when running impala query
- [HUE-4164](#) - The ApiHelper should treat any negative status in the response as an error
- [HUE-4177](#) - Horizontal scroll in FF (Chrome fine) with touch pad is extremely slow
- [HUE-4201](#) - Add warning about max limit of cells before truncation in the export / download query result
- [HUE-4202](#) - Enable offset param for fetching jobbrowser logs
- [HUE-4215](#) - Reset API_CACHE on logout
- [HUE-4224](#) - 'Did you know' on home page is gone
- [HUE-4227](#) - Fix unittest for MR API Cache
- [HUE-4238](#) - Ignore history docs in find_jobs_with_no_doc during sync documents
- [HUE-4238](#) - Ignore history docs in find_jobs_with_no_doc during sync documents
- [HUE-4252](#) - Handle 307 redirect from YARN upon standby failover
- [HUE-4252](#) - Handle 307 redirect from YARN upon standby failover
- [HUE-4253](#) - Prompt for variables just once per variable name
- [HUE-4258](#) - Close and pool Spark History Server connections
- [HUE-4265](#) - Bring back the show preview in the assist
- [HUE-4300](#) - Avoid double file listing call on folder search
- [HUE-4321](#) - Batch submit of SQL show USE the correct DB
- [HUE-4333](#) - Properly reset API_CACHE on failover
- [HUE-4346](#) - Query History disappeared after upgrade to 3.10
- [HUE-4353](#) - Typing in the search bar always redirect to the end of the input
- [HUE-4362](#) - List more oozie workflow parameters on the workflow dashboard page
- [HUE-4364](#) - Handle files with carriage return in create table from a file
- [HUE-4365](#) - No information surfaced when LOAD data from Create table from file fails
- [HUE-4375](#) - Horizontal scrollbar can be hidden under the first fixed column
- [HUE-4383](#) - Trashed queries are showing up in the list of saved queries
- [HUE-4406](#) - Fails to start if Hive/Impala Not Installed
- [HUE-4409](#) - Main right scrollbar does not scoll when on the very right of the screen
- [HUE-4411](#) - Enable scrolling past the end of the editor
- [HUE-4412](#) - Errors should scroll to the line AND the column too
- [HUE-4477](#) - Select All is not filtering out the non visible roles from the selection
- [HUE-4493](#) - Fix sync-workflow action when Workflow includes sub-workflow
- [HUE-4515](#) - Remove oozie.bundle.application.path from properties when rerunning workflow
- [HUE-4533](#) - Disable password reveal on IE
- [HUE-4537](#) - Fix database_logging in hue config so it logs debug database messages
- [HUE-4541](#) - fixing Hue job browser - Kerberos mutual authentication error in Hue
- [HUE-4564](#) - Log stderr on failure to coerce password from script
- [HUE-4616](#) - Only select the snippet DB when executing the first statement
- [HUE-4635](#) - Fix duration on jobs page for running jobs
- [HUE-4662](#) - fixing Hue - Wildcard Certificates not supported
- [HUE-4700](#) - Protect against setting XSS in old editor
- [HUE-4738](#) - Use Concurrency and Throttle values set in coordinator settings
- [HUE-4739](#) - fixed Jobbrowser tests which were failing after resource manager pool change
- [HUE-4766](#) - Replace illegal characters on CSV downloads

- [HUE-4781](#) - Fix export to hdfs to use `download_cell_limit` from `beeswax.conf`
- [HUE-4801](#) - When importing oozie documents and remapping UUIDs, data should be updated accordingly
- [HUE-4808](#) - Don't show the edit link for sub workflows when submitted outside Hue
- [IMPALA-1346](#) - /1590/2344: fix sorter buffer mgmt when spilling
- [IMPALA-3159](#) - impala-shell does not accept wildcard or SAN certificates
- [IMPALA-3344](#) - Simplify sorter and document/enforce invariants.
- [IMPALA-3441](#) - , IMPALA-3659: check for malformed Avro data
- [IMPALA-3499](#) - Split catalog update.
- [IMPALA-3628](#) - Fix cancellation from shell when security is enabled
- [IMPALA-3633](#) - cancel fragment if coordinator is gone
- [IMPALA-3646](#) - Handle corrupt RLE literal or repeat counts of 0.
- [IMPALA-3670](#) - fix sorter buffer mgmt bugs
- [IMPALA-3678](#) - Fix migration of predicates into union operands with an order by + limit.
- [IMPALA-3680](#) - Cleanup the scan range state after failed hdfs cache reads
- [IMPALA-3711](#) - Remove unnecessary privilege checks in `getDbMetadata()`.
- [IMPALA-3732](#) - handle string length overflow in avro files
- [IMPALA-3745](#) - parquet invalid data handling
- [IMPALA-3754](#) - fix `TestParquet.test_corrupt_rle_counts` flakiness
- [IMPALA-3772](#) - Fix admission control stress test.
- [IMPALA-3776](#) - fix 'describe formatted' for Avro tables
- [IMPALA-3820](#) - Handle linkage errors while loading Java UDFs in Catalog
- [IMPALA-3861](#) - Replace `BetweenPredicates` with their equivalent `CompoundPredicate`.
- [IMPALA-3915](#) - Register privilege and audit requests when analyzing resolved table refs.
- [IMPALA-3930](#) - Fix shuffle insert hint with constant partition exprs.
- [IMPALA-3940](#) - Fix getting column stats through views.
- [IMPALA-3965](#) - `TSSLSocketWithWildcardSAN.py` not exported as part of `impala-shell` build lib
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4049](#) - fix empty batch handling NLJ build side
- [OOZIE-2068](#) - Configuration as part of `sharelib`
- [OOZIE-2314](#) - Unable to kill old instance child job by workflow or coord rerun by Launcher
- [OOZIE-2329](#) - Make handling yarn restarts configurable
- [OOZIE-2345](#) - Parallel job submission for forked actions
- [OOZIE-2347](#) - AmendRemove unnecessary new `Configuration()/new jobConf()` calls from oozie
- [OOZIE-2347](#) - amendments patch toRemove unnecessary new `Configuration()/new jobConf()` calls from oozie
- [OOZIE-2347](#) - Remove unnecessary new `Configuration()/new jobConf()` calls from oozie
- [OOZIE-2436](#) - Fork/join workflow fails with `oozie.action.yarn.tag` must not be null
- [OOZIE-2504](#) - Create a `log4j.properties` under `HADOOP_CONF_DIR` in Shell Action
- [OOZIE-2533](#) - Patch-1550 - workaround for
- [OOZIE-2555](#) - Oozie SSL enable setup does not return port for admin -servers
- [OOZIE-2567](#) - HCat connection is not closed while getting hcat cred
- [OOZIE-2589](#) - `CompletedActionXCommand` is hardcoded to wrong priority
- [OOZIE-2649](#) - Can't override sub-workflow configuration property if defined in parent workflow XML
- [OOZIE-2656](#) - `OozieShareLibCLI` uses op system username instead of Kerberos to upload jars
- [PIG-3807](#) - Pig creates wrong schema after dereferencing nested tuple fields with sorts
- [SENTRY-1201](#) - Sentry ignores database prefix for MSCK statement
- [SENTRY-1311](#) - Improve usability of URI privileges by supporting mixed use of URIs with and without scheme
- [SENTRY-1320](#) - Queries of the form `TRUNCATE TABLE db_name.table_name;` no longer fail. The precondition checks allow two child nodes
- [SENTRY-1345](#) - Revert "ACLs on table folder disappear after insert for unpartitioned tables (Sravva Tirukkavalur, Reviewed by: Hao Hao and Anne Yu)"

- [SENTRY-1345](#) - ACLS on table folder disappear after insert for unpartitioned tables
- [SOLR-6295](#) - Fix child filter query creation to never match parent docs in SolrExampleTests
- [SOLR-7280](#) - Missing test resources
- [SOLR-7280](#) - BackportLoad cores in sorted order and tweak coreLoadThread counts to improve cluster stability on restarts
- [SOLR-7866](#) - Harden code to prevent an unhandled NPE when trying to determine the max value of the version field.
- [SOLR-9091](#) - ZkController#publishAndWaitForDownStates logic is inefficient
- [SOLR-9236](#) - AutoAddReplicas will append an extra /tlog to the update log location on replica failover.
- [SPARK-8428](#) - Fix integer overflows in TimSort
- [SPARK-12339](#) - Added a null check that was removed in
- [SPARK-13242](#) - codegen fallback in case-when if there many branches
- [SPARK-14391](#) - Fix launcher communication test, take 2.
- [SPARK-14963](#) - Fix typo in YarnShuffleService recovery file name
- [SPARK-14963](#) - Using recoveryPath if NM recovery is enabled
- [SPARK-15165](#) - Introduce place holder for comments in generated code
- [SPARK-16106](#) - TaskSchedulerImpl should properly track executors added to existing hosts
- [SPARK-16505](#) - Optionally propagate error during shuffle service startup.
- [SQOOP-2561](#) - Special Character removal from Column name as avro data results in duplicate column and fails the import
- [SQOOP-2846](#) - Sqoop Export with update-key failing for avro data file
- [SQOOP-2906](#) - Optimization of AvroUtil.toAvroIdentifier
- [SQOOP-2920](#) - sqoop performance deteriorates significantly on wide datasets; sqoop 100% on cpu
- [SQOOP-2971](#) - OraOop does not close connections properly
- [SQOOP-2995](#) - Backward incompatibility introduced by Custom Tool options.
- [SQOOP-2999](#) - Sqoop ClassNotFoundException (org.apache.commons.lang3.StringUtils) is thrown when executing Oracle direct import map task

Issues Fixed in CDH 5.8.0

CDH 5.8.0 fixes the following issues.

Apache Flume

Flume fully compatible with Kafka 2.x

In release CDH 5.8.0, Flume is fully compatible with Kafka 2.x, including support for security features.

Apache HBase

Premature EOF detected in a WAL During Replication

Cloudera Bug: CDH-38113

During the parsing of a write-ahead log (WAL) during replication, an `InvalidProtobufException` can occur while reading the source RegionServer WAL, if `EOF` (end-of-file) is incorrectly detected before the actual end of the file. HBase stops reading the WAL after the EOF, and does not parse any bytes which occur after the EOF, causing data loss.

To work around this problem, Cloudera has patched HBase. HBase in CDH 5.8.0 and higher detect whether unparsed bytes exist after the `EOF`, and if so, the WAL is reset and re-read from the beginning, to attempt a clean read-through.

In testing, a single reset has been sufficient to work around observed data loss. However, the above change will retry a given WAL file indefinitely. On each attempt, a log message such as this will be emitted at the `WARN` level:

```
Processing end of WAL file '{}'. At position {}, which is too far away from
reported file length {}. Restarting WAL reading
```

Additional log detail are emitted at the `TRACE` level about file offsets seen while handling recoverable errors.

*Batch Get after Batch Put Does Not Fetch All Cells***Bug:** [HBASE-15811](#)**Cloudera Bug:** CDH-40344

A batch Get after a batch Put could fail to fetch cells that were written by the Get, resulting in a "read-your-writes" failure. This bug was exacerbated by high load on the client.

*Read Replica Failure For PUT Operation During Region Transition***Cloudera Bug:** CDH-39758

When the patch for HBASE-10794 was applied in CDH 5.4.4, a new bug was introduced, where, if the primary RegionServer becomes unavailable (for any reason, even a graceful shutdown), while a client is performing PUTs on that region, subsequent PUTs will fail.

*Latency Metrics Inaccurate for MultiGet Operations***Bug:** [HBASE-15673](#)**Cloudera Bug:** CDH-39422

Latency values are written after each row is processed. However, if `MultiGet` is enabled, some rows are not counted in the metrics. This causes the metrics for the 50th, 75th, and 90th percentiles to be reported as 0.

*Inconsistent Behavior Among DeleteColumnFamilyProcedure, CreateTableProcedure, and ModifyTableProcedure***Bug:** [HBASE-15456](#)**Cloudera Bug:** CDH-12345

If there is only one family in the table, `DeleteColumnFamilyProcedure` will fail. When `hbase.table.sanity.checks` is set to `false`, the HMaster logs a warning, but `CreateTableProcedure` and `ModifyTableProcedure` will now fail, where before they logged a warning, but succeeded. This makes the behavior of all three methods consistent.

*Failed hbase-spark Bulk Loads Leave Files Behind***Bug:** [HBASE-15271](#)**Cloudera Bug:** CDH-38145

When using the bulk load helper provided by the `hbase-spark` module, output files are now written into temporary files and only made available when the executor has successfully completed. Previously, failed executors would leave files behind, and these files would be picked up by subsequent bulk load commands, and spurious copies of some cells were written.

Apache Hive

[HIVE-13217](#), CDH-30121: Replication for HoS MapJoin small file needs to respect `dfs.replication.max`[HIVE-13039](#), CDH-37322: BETWEEN predicate is not functioning correctly with predicate PUSHDOWN on Parquet table[HIVE-13065](#), CDH-37409: Hive throws `NullPointerException` (NPE) when writing map type data to an HBase-backed table[HIVE-13160](#), CDH-37847: HS2 unable to load UDFs on startup when HMS is not ready[HIVE-13243](#), CDH-38477: Hive DROP TABLE on encryption zone fails for external tables[HIVE-13302](#), CDH-38581: Direct SQL: CAST to DATE doesn't work on Oracle[HIVE-13115](#), CDH-38612: MetaStore Direct SQL `getPartitions()` call fails when the columns schemas for a partition are NULL[HIVE-10303](#), CDH-38685: HIVE-9471 broke forward compatibility of ORC files

- [HIVE-12706](#), CDH-39108: Incorrect output from `from_utc_timestamp()` / `to_utc_timestamp` when local timezone has DST
- [HIVE-10685](#), CDH-39187: ALTER TABLE concatenate operator will cause duplicate data
- [HIVE-13500](#): Launching big queries fails with `OutOfMemoryException`
- [HIVE-13527](#), CDH-39616: Using deprecated APIs in HBase client causes ZooKeeper connection leaks.
- [HIVE-12517](#), CDH-39722: Beeline's use of failed connection(s) causes failures and leaks.
- [HIVE-13632](#), CDH-39911: Hive failing on INSERT empty array into parquet table
- [HIVE-13285](#), CDH-39951: Orc concatenation may drop old files from moving to final path
- [HIVE-13836](#), CDH-40070: `DbNotifications` giving an error = Invalid state. Transaction has already started
- [HIVE-9499](#), CDH-40478: `hive.limit.query.max.table.partition` makes queries fail on non-partitioned tables
- [HIVE-13462](#), CDH-41031: `HiveResultSetMetaData.getPrecision()` fails for NULL columns
- [HIVE-11408](#), CDH-32685: HiveServer2 is leaking ClassLoaders when add jar / temporary functions are used due to constructor caching in Hadoop `ReflectionUtils`
- [HIVE-12481](#), CDH-36124: Occasionally "Request is a replay" will be thrown from HS2
- [HIVE-10698](#), CDH-37320: Query on view results fails with "table not found error" if view is created with subquery alias (CTE)
- [HIVE-12941](#), CDH-37979: Unexpected result when using `MIN()` on struct with NULL in first field
- [HIVE-13200](#), CDH-38046: Aggregation functions returning empty rows on partitioned columns
- [HIVE-11054](#), CDH-38517: Read error : Partition Varchar column cannot be cast to string
- [HIVE-13401](#), CDH-39009: Kerberized HS2 with LDAP auth enabled fails kerberos/delegation token authentication
- [HIVE-13217](#): Some queries with UNION all fail when CBO is off
- [HIVE-11369](#), CDH-39715: MapJoins in HiveServer2 fail when `jmxremote` is used
- [HIVE-13261](#), CDH-40042: Can not compute column stats for partition when schema evolves

With Sentry enabled, only Hive admin users have access to YARN job logs

As a prerequisite of enabling Sentry, Hive impersonation is turned off, which means all YARN jobs are submitted to the Hive job queue, and are run as the `hive` user. This is an issue because the YARN History Server now has to block users from accessing logs for their own jobs, since their own usernames are not associated with the jobs. As a result, end users cannot access any job logs unless they can get `sudo` access to the cluster as the `hdfs`, `hive` or other admin users.

In CDH 5.8 (and higher), Hive overrides the default configuration, `mapred.job.queue.name`, and places incoming jobs into the connected user's job queue, even though the submitting user remains `hive`. Hive obtains the relevant queue/username information for each job by using YARN's `fair-scheduler.xml` file.

Hue

Cannot query the `customers` table in Hue

Bug: [HUE-3040](#)

Cloudera Bug: CDH-33974

To query the `customers` table, users must re-create the parquet data for compatibility.

Cloudera Distribution of Apache Kafka

CDH 5.7 is not compatible with Cloudera Distribution of Apache Kafka 1.x

Cloudera Distribution of Apache Kafka 1.x is compatible with CDH 5.4+.

Apache Oozie

PySpark does not work from the Oozie Spark Action

Bug: [OOZIE-2482](#)

Cloudera Bug: CDH-36349

The Spark Action would typically fail with a message like, "key not found: SPARK_HOME," but other error messages were possible. After the fix, the Spark Action has the necessary changes to successfully run PySpark jobs. See [Oozie Spark Action Extension](#) for more details and an example. Cloudera makes the PySpark dependencies available.

Apache Sentry

Security

Sentry does not check privileges on the URI used for the `CREATE INDEX LOCATION '/path'` command

Bug: [SENTRY-1231](#)

The `CREATE INDEX LOCATION '/path'` command would succeed even if a user did not have the required URI privileges for the `/path`.

Upgraded `libthrift` to version 0.9.3 due to a security vulnerability

Cloudera Bug: CDH-40034

For details on the security vulnerability in the Apache Thrift client libraries, see [THRIFT-3231](#).

Hive Binding

`INSERT OVERWRITE DIRECTORY` command does not work correctly

Bug: [SENTRY-922](#)

The `INSERT OVERWRITE DIRECTORY` command would write table data into an HDFS directory (`hdfs://path/`), even if privileges are granted only for the local directory (`file://path/`).

`INSERT INTO` no longer requires URI privilege on partition locations

Bug: [SENTRY-1095](#)

The `INSERT INTO` Hive command adds location information to the partition description. Usually if location information is included, you must ensure that the user has privileges on the corresponding URI. However, in this case, since the partition locations are under the table directory and can be easily generated, these requirements have been relaxed.

Change default value of `sentry.hive.server`

Bug: [SENTRY-1112](#)

The default value for `sentry.hive.server` was changed from `server1` to an empty string.

Sentry Service

Sentry's Oracle upgrade scripts fails with `ORA-00955`

Bug: [SENTRY-1066](#)

Sentry upgrade scripts for Oracle would fail with error, `ORA-00955`, because during the upgrade, the script inadvertently creates an index with the same name as the constraint being dropped. The script will now run `DROP INDEX` before it adds the constraint again and completes the schema upgrade successfully.

`grantServerPrivilege()` and `revokeServerPrivilege()` should treat '*' and 'ALL' as synonyms

Bug: [SENTRY-1252](#)

The `grantServerPrivilege()` and `revokeServerPrivilege()` methods should treat `*` and `ALL` as synonyms when an action is not explicitly specified. Previously, if `grantServerPrivilege()` was called without an action, and followed up with a `revokeServerPrivilege()` invocation with an action such as `ALL`, the server-level privilege would not be revoked. This fix only applies to privileges that are granted after upgrading to CDH 5.8.

Sentry Debugging

Error in Hive Metastore Plugin (`renameAuthzObject`) log messages

Bug: [SENTRY-1169](#)

The `renameAuthzObject` plugin prints log messages with old path names in place of new path names.

Apache ZooKeeper

Upgrade Netty Due to Security Vulnerabilities

Bug: [ZOOKEEPER-2450](#)

Cloudera Bug: CDH-39988

Netty was upgraded from version 3.2.2 to 3.10.5 to resolve security vulnerabilities.

Fix Privacy Violation in Login.java

Bug: [ZOOKEEPER-2405](#)

Cloudera Bug: CDH-36673

In `Login.java`, `getTGT()` was logging confidential information in `DEBUG` mode. After the fix, only principals are logged.

Issues Fixed in CDH 5.7.x

The following topics describe issues fixed in CDH 5.7.x, from newest to oldest release. You can also review [What's New In CDH 5.7.x](#) on page 1829 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.7.6

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.7.6:

- [AVRO-1943](#) - Flaky test: `TestNettyServerWithCompression.testConnectionsCount`.
- [FLUME-2908](#) - `NetcatSource` - `SocketChannel` not closed when session is broken.
- [FLUME-2997](#) - Fix flaky test in `SpillableMemoryChannel`.
- [FLUME-3002](#) - Fix tests in `TestBucketWriter`.
- [FLUME-3003](#) - Fix flaky test `SourceCounter` in `TestSyslogUdpSource`.
- [FLUME-3049](#) - Make HDFS sink rotate more reliably in secure mode.
- [HADOOP-7930](#) - Kerberos relogin interval in `UserGroupInformation` should be configurable.
- [HADOOP-11031](#) - Design Document for Credential Provider API.
- [HADOOP-11619](#) - `FTPFileSystem` should override `getDefaultPort`.
- [HADOOP-12453](#) - Support decoding KMS Delegation Token with its own Identifier.
- [HADOOP-12537](#) - S3A to support Amazon STS temporary credentials.
- [HADOOP-12655](#) - `TestHttpServer.testBindAddress` bind port range is wider than expected.
- [HADOOP-12723](#) - S3A: Add ability to plug in any `AWSCredentialsProvider`.
- [HADOOP-13034](#) - Log message about input options in `distcp` lacks some items.
- [HADOOP-13433](#) - Race in `UGI.reloginFromKeytab`.
- [HADOOP-13590](#) - Retry until TGT expires even if the UGI renewal thread encountered exception.
- [HADOOP-13627](#) - Have an explicit `KerberosAuthException` for UGI to throw, text from public constants.
- [HADOOP-13641](#) - Update `UGI#spawnAutoRenewalThreadForUserCreds` to reduce indentation.
- [HADOOP-13838](#) - `KMSTokenRenewer` should close providers.
- [HADOOP-13953](#) - Make `FTPFileSystem`'s data connection mode and transfer mode configurable.

- [HADOOP-14003](#) - Make additional KMS tomcat settings configurable.
- [HDFS-9428](#) - Fix intermittent failure of TestDNFencing.testQueueingWithAppend.
- [HDFS-9630](#) - DistCp minor refactoring and clean up.
- [HDFS-9638](#) - Improve DistCp Help and documentation.
- [HDFS-9764](#) - DistCp doesn't print value for several arguments including -numListstatusThreads.
- [HDFS-9804](#) - Allow long-running Balancer to login with keytab.
- [HDFS-9820](#) - Improve distcp to support efficient restore to an earlier snapshot.
- [HDFS-9888](#) - Allow resetting KerberosName in unit tests.
- [HDFS-10216](#) - Distcp -diff throws exception when handling relative path.
- [HDFS-10271](#) - Extra bytes are getting released from reservedSpace for append.
- [HDFS-10298](#) - Document the usage of distcp -diff option.
- [HDFS-10313](#) - Distcp need to enforce the order of snapshot names passed to -diff.
- [HDFS-10336](#) - TestBalancer failing intermittently because of not resetting UserGroupInformation completely
- [HDFS-10397](#) - Distcp should ignore -delete option if -diff option is provided instead of exiting.
- [HDFS-10556](#) - DistCpOptions should be validated automatically.
- [HDFS-10763](#) - Open files can leak permanently due to inconsistent lease update.
- [HDFS-11040](#) - Add documentation for HDFS-9820 distcp improvement.
- [HDFS-11056](#) - Concurrent append and read operations lead to checksum error.
- [HDFS-11160](#) - VolumeScanner reports write-in-progress replicas as corrupt incorrectly.
- [HDFS-11229](#) - HDFS-11056 failed to close meta file.
- [HDFS-11275](#) - Check groupEntryIndex and throw a helpful exception on failures when removing ACL.
- [HDFS-11292](#) - log lastWrittenTxId etc info in logSyncAll.
- [HDFS-11306](#) - Print remaining edit logs from buffer if edit log can't be rolled.
- [MAPREDUCE-6571](#) - JobEndNotification info logs are missing in AM container syslog.
- [MAPREDUCE-6763](#) - Shuffle server listen queue is too small.
- [MAPREDUCE-6798](#) - Fix intermittent failure of TestJobHistoryParsing.testJobHistoryMethods.
- [MAPREDUCE-6801](#) - Fix flaky TestKill.testKillJob.
- [MAPREDUCE-6817](#) - The format of job start time in JHS is different from those of submit and finish time.
- [MAPREDUCE-6831](#) - Flaky test TestJobImpl.testKilledDuringKillAbort
- [YARN-2306](#) - Add test for leakage of reservation metrics in fair scheduler.
- [YARN-3554](#) - Default value for maximum nodemanager connect wait time is too high.
- [YARN-4363](#) - In TestFairScheduler, testcase should not create FairScheduler redundantly.
- [YARN-4555](#) - TestDefaultContainerExecutor#testContainerLaunchError fails on non-english locale environment.
- [YARN-5752](#) - TestLocalResourcesTrackerImpl#testLocalResourceCache times out.
- [YARN-5837](#) - NPE when getting node status of a decommissioned node after an RM restart.
- [YARN-5859](#) - TestResourceLocalizationService#testParallelDownloadAttemptsForPublicResource sometimes fails
- [YARN-5862](#) - TestDiskFailures.testLocalDirsFailures failed.
- [YARN-5890](#) - FairScheduler should log information about AM-resource-usage and max-AM-share for queues.
- [YARN-5920](#) - Fix deadlock in TestRMHA.testTransitionedToStandbyShouldNotHang.
- [HBASE-15324](#) - Jitter may cause desiredMaxFileSize overflow in ConstantSizeRegionSplitPolicy and trigger unexpected split.
- [HBASE-15430](#) - Failed taking snapshot - Manifest proto-message too large.
- [HBASE-16172](#) - Unify the retry logic in ScannerCallableWithReplicas and RpcRetryingCallerWithReadReplicas.
- [HBASE-16270](#) - Handle duplicate clearing of snapshot in region replicas.
- [HBASE-16345](#) - RpcRetryingCallerWithReadReplicas#call() should catch some RegionServer Exceptions.
- [HBASE-16824](#) - Writer.flush() can be called on already closed streams in WAL roll.
- [HBASE-16841](#) - Data loss in MOB files after cloning a snapshot and deleting that snapshot.
- [HBASE-17058](#) - Lower epsilon used for jitter verification from HBASE-15324.
- [HBASE-17241](#) - Avoid compacting already compacted mob files with _del files.
- [HBASE-17452](#) - Failed taking snapshot - region Manifest proto-message too large.

- [HBASE-17522](#) - Handle JVM throwing runtime exceptions when we ask for details on heap usage the same as a correctly returned 'undefined'.
- [HIVE-10965](#) - direct SQL for stats fails in 0-column case.
- [HIVE-11849](#) - NPE in HiveHBaseTableShapshotInputFormat in query with just count(*).
- [HIVE-12083](#) - HIVE-10965 introduces thrift error if partNames or colNames are empty.
- [HIVE-12465](#) - Hive might produce wrong results when (outer) joins are merged.
- [HIVE-12619](#) - Switching the field order within an array of structs causes the query to fail.
- [HIVE-12780](#) - Fix the output of the history command in Beeline HIVE-12789: Fix output twice in the history command of Beeline.
- [HIVE-12891](#) - Hive fails when java.io.tmpdir is set to a relative location.
- [HIVE-12976](#) - MetaStoreDirectSql doesn't batch IN lists in all cases.
- [HIVE-13129](#) - CliService leaks HMS connection.
- [HIVE-13149](#) - Remove some unnecessary HMS connections from HS2.
- [HIVE-13240](#) - GroupByOperator: Drop the hash aggregates when closing operator.
- [HIVE-13539](#) - HiveHFileOutputFormat searching the wrong directory for HFiles.
- [HIVE-13696](#) - Modify FairSchedulerShim to dynamically reload changes to fair-scheduler.xml.
- [HIVE-13866](#) - flatten callstack for directSQL errors.
- [HIVE-14173](#) - NPE was thrown after enabling directsql in the middle of session.
- [HIVE-14764](#) - Enabling "hive.metastore.metrics.enabled" throws OOM in HiveMetastore.
- [HIVE-14820](#) - RPC server for spark inside HS2 is not getting server address properly.
- [HIVE-15054](#) - Hive insertion query execution fails on Hive on Spark.
- [HIVE-15090](#) - Temporary DB failure can stop ExpiredTokenRemover thread.
- [HIVE-15338](#) - Wrong result from non-vectorized DATEDIFF with scalar parameter of type DATE/TIMESTAMP.
- [HIVE-15410](#) - WebHCat supports get/set table property with its name containing period and hyphen.
- [HIVE-15551](#) - memory leak in directsql for mysql+bonecp specific initialization.
- [HUE-4662](#) - [security] fixing Hue - Wildcard Certificates not supported.
- [HUE-4466](#) - [security] deliver csrftoken cookie with secure bit set if possible.
- [HUE-5163](#) - [security] Speed up initial page rendering.
- [HUE-4916](#) - [core] Truncate last name to 30 chars on ldap import.
- [HUE-5050](#) - [core] Logout fails for local login when multiple backends are used.
- [HUE-5042](#) - [core.backend] Unable to kill jobs after Resource Manager failover.
- [HUE-4201](#) - [editor] Add warning about max limit of cells before truncation.
- [HUE-4968](#) - [oozie] Remove access to /oozie/import_workflow when v2 is enabled.
- [HUE-5218](#) - [search] Validate dashboard sharing works.
- [IMPALA-2864](#) - Ensure that client connections are closed after a failed Open()
- [IMPALA-3167](#) - Fix assignment of WHERE conjunct through grouping agg + OJ.
- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3698](#) - Fix Isilon permissions test
- [IMPALA-3861](#) - Replace BetweenPredicates with their equivalent CompoundPredicate.
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3983](#) - /IMPALA-3974: Delete function jar resources after load
- [IMPALA-4153](#) - Return valid non-NULL pointer for 0-byte allocations
- [IMPALA-4223](#) - Handle truncated file read from HDFS cache
- [IMPALA-4336](#) - Cast exprs after unnesting union operands.
- [IMPALA-4363](#) - Add Parquet timestamp validation
- [IMPALA-4423](#) - Correct but conservative implementation of Subquery.equals().
- [IMPALA-4433](#) - Always generate testdata using the same time zone setting
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog
- [IMPALA-4488](#) - HS2 GetOperationStatus() should keep session alive
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots

- [IMPALA-4579](#) - SHOW CREATE VIEW fails for view containing a subquery
- [IMPALA-4765](#) - Avoid using several loading threads on one table.
- [OOZIE-2194](#) - oozie job -kill doesn't work with spark action.
- [OOZIE-2243](#) - Kill Command does not kill the child job for java action.
- [OOZIE-2584](#) - Eliminate Thread.sleep() calls in TestMemoryLocks.
- [OOZIE-2678](#) - Oozie job -kill doesn't work with tez jobs.
- [OOZIE-2742](#) - Unable to kill applications based on tag.
- [PIG-5025](#) - Fix flaky test failures in TestLoad.java.
- [SENTRY-1260](#) - Improve error handling - ArrayIndexOutOfBoundsException in PathsUpdate.parsePath can cause MetastoreCacheInitializer initialization to fail.
- [SENTRY-1270](#) - Improve error handling - Database with malformed URI causes NPE in HMS plugin during DDL.
- [SENTRY-1520](#) - Provide mechanism for triggering HMS full snapshot.
- [SENTRY-1564](#) - Improve error detection and reporting in MetastoreCacheInitializer.java.
- [SOLR-9284](#) - The HDFS BlockDirectoryCache should not let it's keysToRelease or names maps grow indefinitely.
- [SOLR-9330](#) - Fix AlreadyClosedException on admin/mbeans?stats=true.
- [SOLR-10031](#) - Validation of filename params in ReplicationHandler.
- [SPARK-12241](#) - [YARN] Improve failure reporting in Yarn client obtainTokenForHBase().
- [SPARK-12523](#) - [YARN] Support long-running of the Spark On HBase and hive meta store.
- [SPARK-12966](#) - [SQL] ArrayType(DecimalType) support in Postgres JDBC.
- [SPARK-13566](#) - [CORE] Avoid deadlock between BlockManager and Executor Thread.
- [SPARK-13958](#) - Executor OOM due to unbounded growth of pointer array in...
- [SPARK-14204](#) - [SQL] register driverClass rather than user-specified class.
- [SPARK-16044](#) - [SQL] Backport input_file_name() for data source based on NewHadoopRDD to branch 1.6.
- [SPARK-17245](#) - [SQL][BRANCH-1.6] Do not rely on Hive's session state to retrieve HiveConf.
- [SPARK-17465](#) - [SPARK CORE] Inappropriate memory management in `org.apache.spark.storage.MemoryStore` may lead to memory leak.
- [SPARK-18750](#) - [YARN] Follow up: move test to correct directory in 2.1 branch.
- [SPARK-18750](#) - [YARN] Avoid using "mapValues" when allocating containers.
- [SQOOP-2349](#) - Add command line option for setting transaction isolation levels for metadata queries.
- [SQOOP-2880](#) - Provide argument for overriding temporary directory.
- [SQOOP-2884](#) - Document --temporary-rootdir.
- [SQOOP-2915](#) - Fixing Oracle related unit tests.
- [SQOOP-2983](#) - OraOop export has degraded performance with wide tables.
- [SQOOP-3013](#) - Configuration "tmpjars" is not checked for empty strings before passing to MR.
- [SQOOP-3028](#) - Include stack trace in the logging of exceptions in ExportTool.
- [SQOOP-3034](#) - HBase import should fail fast if using anything other than as-textfile.
- [SQOOP-3053](#) - Create a cmd line argument for sqoop.throwOnError and use it through SqoopOptions.
- [SQOOP-3055](#) - Fixing MySQL tests failing due to ignored test inputs/configuration.
- [SQOOP-3057](#) - Fixing 3rd party Oracle tests failing due to invalid case of column names.
- [SQOOP-3071](#) - Fix OracleManager to apply localTimeZone correctly in case of Date objects too.
- [SQOOP-3124](#) - Fix ordering in column list query of PostgreSQL connector to reflect the logical order instead of adhoc ordering.

Issues Fixed in CDH 5.7.5

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.7.5:

- [HADOOP-10300](#) - Allowed deferred sending of call responses
- [HADOOP-12483](#) - Maintain wrapped SASL ordering for postponed IPC responses
- [HADOOP-13317](#) - Add logs to KMS server-side to improve supportability
- [HADOOP-13558](#) - UserGroupInformation created from a Subject incorrectly tries to renew the Kerberos ticket

- [HADOOP-13638](#) - KMS should set UGI's Configuration object properly
- [HADOOP-13669](#) - KMS Server should log exceptions before throwing
- [HADOOP-13693](#) - Remove the message about HTTP OPTIONS in SPNEGO initialization message from kms audit log
- [HDFS-4176](#) - EditLogTailer should call rollEdits with a timeout
- [HDFS-6962](#) - ACLs inheritance conflict with umaskmode
- [HDFS-7413](#) - Some unit tests should use NameNodeProtocols instead of FSNameSystem
- [HDFS-7964](#) - Add support for async edit logging
- [HDFS-8224](#) - Schedule a block for scanning if its metadata file is corrupt
- [HDFS-8709](#) - Clarify automatic sync in FSEditLog#logEdit
- [HDFS-9038](#) - DFS reserved space is erroneously counted towards non-DFS used
- [HDFS-10178](#) - Permanent write failures can happen if pipeline recoveries occur for the first packet
- [HDFS-10609](#) - Uncaught InvalidEncryptionKeyException during pipeline recovery can abort downstream applications
- [HDFS-10641](#) - TestBlockManager#testBlockReportQueueing fails intermittently
- [HDFS-10722](#) - Fix race condition in TestEditLog#testBatchedSyncWithClosedLogs
- [HDFS-10760](#) - DataXceiver#run() should not log InvalidToken exception as an error
- [HDFS-10879](#) - TestEncryptionZonesWithKMS#testReadWrite fails intermittently
- [HDFS-10962](#) - TestRequestHedgingProxyProvider fails intermittently
- [HDFS-11012](#) - Unnecessary INFO logging on DFSClients for InvalidToken
- [MAPREDUCE-6633](#) - AM should retry map attempts if the reduce task encounters compression related errors
- [MAPREDUCE-6718](#) - Add progress log to JHS during startup
- [MAPREDUCE-6728](#) - Give fetchers hint when ShuffleHandler rejects a shuffling connection
- [MAPREDUCE-6771](#) - RMContainerAllocator sends container diagnostics event after corresponding completion event
- [YARN-4004](#) - container-executor should print output of Docker logs if the Docker container exits with non-0 exit status
- [YARN-4017](#) - container-executor overuses PATH_MAX
- [YARN-4245](#) - Generalize config file handling in container-executor
- [YARN-4255](#) - container-executor does not clean up Docker operation command files
- [YARN-4723](#) - NodesListManager\$UnknownNodeId ClassCastException
- [YARN-4940](#) - YARN node -list -all fails if RM starts with decommissioned node
- [YARN-5704](#) - Provide configuration knobs to control enabling/disabling new/work in progress features in container-executor
- [HBASE-16294](#) - hbck reporting "No HDFS region dir found" for replicas
- [HBASE-16699](#) - Overflows in AverageIntervalRateLimiter's refill() and getWaitInterval()
- [HBASE-16767](#) - Mob compaction needs to clean up files in /hbase/mobdir/.tmp and /hbase/mobdir/.tmp/.bulkload when running into IO exceptions
- [HIVE-10384](#) - BackportRetryingMetaStoreClient does not retry wrapped TTransportExceptions
- [HIVE-12077](#) - MSCK Repair table should fix partitions in batches
- [HIVE-12475](#) - Parquet schema evolution within array<struct<>> does not work
- [HIVE-12785](#) - View with union type and UDF to the struct is broken
- [HIVE-13058](#) - Add session and operation_log directory deletion messages
- [HIVE-13198](#) - Authorization issues with cascading views
- [HIVE-13237](#) - Select parquet struct field with upper case throws NPE
- [HIVE-13429](#) - Tool to remove dangling scratch dir
- [HIVE-13997](#) - Insert overwrite directory does not overwrite existing files
- [HIVE-14313](#) - Test failure TestMetaStoreMetrics.testConnections
- [HIVE-14421](#) - FS.deleteOnExit holds references to _tmp_space.db files
- [HIVE-14762](#) - Add logging while removing scratch space
- [HIVE-14784](#) - Operation logs are disabled automatically if the parent directory does not exist

- [HIVE-14799](#) - Query operations are not thread safe during cancellation
- [HIVE-14805](#) - Subquery inside a view will have the object in the subquery as the direct input
- [HIVE-14810](#) - Fix failing test: TestMetaStoreMetrics.testMetaDataCounts
- [HIVE-14817](#) - Shutdown the SessionManager timeoutChecker thread properly upon shutdown
- [HIVE-14839](#) - Improve the stability of TestSessionManagerMetrics
- [HUE-3860](#) - Fix unittest beeswax.tests.test_hiveserver2_jdbc_url
- [HUE-3905](#) - Reset beeswax.conf params in beeswax.tests.test_hiveserver2_jdbc_url
- [HUE-4201](#) - Add warning about max limit of cells before truncation in the download query result
- [HUE-4662](#) - Fixed: Wildcard Certificates not supported
- [HUE-4739](#) - Fixed Jobbrowser tests which were failing after resource manager pool change
- [HUE-4916](#) - Truncate last name to 30 chars on ldap import
- [HUE-4968](#) - Remove access to /oozie/import_workflow when v2 is enabled
- [HUE-5042](#) - Unable to kill jobs after Resource Manager failover
- [HUE-5050](#) - Logout fails for local login when multiple backends are used
- [HUE-5161](#) - Speed up roles rendering
- [HUE-5163](#) - Speed up initial page rendering
- [IMPALA-1619](#) - Support 64-bit allocations
- [IMPALA-1740](#) - Add support for skip.header.line.count
- [IMPALA-3458](#) - Fix table creation to test insert with header lines
- [IMPALA-3949](#) - Log the error message in FileSystemUtil.copyToLocal()
- [IMPALA-4037](#) - Fix locking during query cancellation
- [IMPALA-4076](#) - Fix runtime filter sort compare method
- [IMPALA-4099](#) - Fix the error message while loading UDFs with no JARs
- [IMPALA-4120](#) - Incorrect results with LEAD() analytic function
- [IMPALA-4135](#) - Thrift threaded server times-out connections during high load
- [IMPALA-4170](#) - Fix identifier quoting in COMPUTE INCREMENTAL STATS
- [IMPALA-4196](#) - Cross compile bit-byte functions
- [IMPALA-4237](#) - Fix materialization of 4 byte decimals in data source scan node
- [IMPALA-4246](#) - SleepForMs() utility function has undefined behavior for > 1s
- [OOZIE-1814](#) - Oozie should mask any passwords in logs and REST interfaces
- [SOLR-9310](#) - PeerSync fails on a node restart due to IndexFingerPrint mismatch
- [SPARK-12009](#) - Avoid re-allocating YARN container when driver wants to stop all Executors
- [SPARK-12392](#) - Optimize a location order of broadcast blocks by considering preferred local hosts
- [SPARK-12941](#) - Spark-SQL JDBC Oracle dialect fails to map string datatypes to Oracle VARCHAR datatype mapping
- [SPARK-13328](#) - Poor read performance for broadcast variables with dynamic resource allocation
- [SPARK-16625](#) - General data types to be mapped to Oracle
- [SPARK-16711](#) - YarnShuffleService does not re-init properly on YARN rolling upgrade
- [SPARK-17171](#) - DAG lists all partitions in the graph
- [SPARK-17433](#) - YarnShuffleService does not handle moving credentials levelDb
- [SPARK-17611](#) - Make shuffle service test really test authentication
- [SPARK-17644](#) - Do not add failedStages when abortStage for fetch failure
- [SPARK-17696](#) - Partial backport of to branch-1.6.
- [SQOOP-2952](#) - Row key not added into column family using --hbase-bulkload
- [SQOOP-2986](#) - Add validation check for --hive-import and --incremental lastmodified
- [SQOOP-3021](#) - ClassWriter fails if a column name contains a backslash character

Issues Fixed in CDH 5.7.4

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.7.4:

- [FLUME-2797](#) - Use SourceCounter for SyslogTcpSource

- [FLUME-2844](#) - SpillableMemoryChannel must start ChannelCounter
- [HADOOP-8436](#) - NPE In getLocalPathForWrite (path, conf) when the required context item is not configured
- [HADOOP-8437](#) - getLocalPathForWrite should throw IOException for invalid paths
- [HADOOP-10048](#) - LocalDirAllocator should avoid holding locks while accessing the filesystem
- [HADOOP-11469](#) - KMS should skip default.key.acl and whitelist.key.acl when loading key acl.
- [HADOOP-12252](#) - LocalDirAllocator should not throw NPE with empty string configuration
- [HADOOP-12548](#) - Read s3a credentials from a Credential Provider
- [HADOOP-12609](#) - Fix intermittent failure of TestDecayRpcScheduler.
- [HADOOP-13270](#) - BZip2CompressionInputStream finds the same compression marker twice in corner case, causing duplicate data blocks
- [HADOOP-13353](#) - LdapGroupsMapping getPassword should not return null when IOException is thrown
- [HADOOP-13437](#) - KMS should reload whitelist and default key ACLs when hot-reloading
- [HADOOP-13487](#) - Hadoop KMS should load old delegation tokens from Zookeeper on startup
- [HADOOP-13526](#) - Add detailed logging in KMS for the authentication failure of proxy user
- [HADOOP-13579](#) - Fix source-level compatibility after HADOOP-11252
- [HDFS-4210](#) - Throw helpful exception when DNS entry for JournalNode cannot be resolved
- [HDFS-7415](#) - Move FSNameSystem.resolvePath() to FSDirectory
- [HDFS-7420](#) - Delegate permission checks to FSDirectory
- [HDFS-7463](#) - Simplify FSNamesystem#getBlockLocationsUpdateTimes
- [HDFS-7478](#) - Move org.apache.hadoop.hdfs.server.namenode.NNConf to FSNamesystem
- [HDFS-7517](#) - Remove redundant non-null checks in FSNamesystem#getBlockLocations
- [HDFS-8269](#) - getBlockLocations() does not resolve the .reserved path and generates incorrect edit logs when updating the atime
- [HDFS-8897](#) - Balancer should handle fs.defaultFS trailing slash in HA
- [HDFS-9198](#) - Coalesce IBR processing in the NameNode.
- [HDFS-9781](#) - FsDatasetImpl#getBlockReports can occasionally throw NullPointerException
- [HDFS-9906](#) - Remove unhelpful log entries when restarting a datanode
- [HDFS-9958](#) - BlockManager#createLocatedBlocks can throw NPE for corruptBlocks on failed storages
- [HDFS-10270](#) - TestJMXGet:testNameNode() fails
- [HDFS-10457](#) - DataNode should not auto-format block pool directory if VERSION is missing
- [HDFS-10544](#) - Balancer does not work with IPFailoverProxyProvider
- [HDFS-10643](#) - Namenode should use loginUser(hdfs) to generateEncryptedKey
- [HDFS-10822](#) - Log DataNodes in the write pipeline
- [MAPREDUCE-4784](#) - TestRecovery occasionally fails
- [MAPREDUCE-6359](#) - In RM HA setup, Cluster tab links populated with AM hostname instead of RM
- [MAPREDUCE-6514](#) - Fixed MapReduce ApplicationMaster to properly updated resources ask after ramping down of all reducers avoiding job hangs
- [MAPREDUCE-6628](#) - Potential memory leak in CryptoOutputStream
- [MAPREDUCE-6670](#) - TestJobListCache#testEviction sometimes fails on Windows with timeout
- [MAPREDUCE-6680](#) - JHS UserLogDir scan algorithm sometimes could skip directory with update in CloudFS (Azure FileSystem, S3, and so on)
- [MAPREDUCE-6684](#) - High contention on scanning of user directory under immediate_done in Job History Server
- [MAPREDUCE-6738](#) - TestJobListCache.testAddExisting failed intermittently in slow VM testbed
- [MAPREDUCE-6761](#) - Regression when handling providers - invalid configuration ServiceConfiguration causes Cluster initialization failure
- [YARN-2977](#) - Fixed intermittent TestNMClient failure
- [YARN-4989](#) - TestWorkPreservingRMRestart#testCapacitySchedulerRecovery fails intermittently
- [YARN-5608](#) - TestAMRMClient.setup() fails with ArrayOutOfBoundsException
- [HBASE-15856](#) - Fix UnknownHostException import in MetaTableLocator
- [HBASE-15856](#) - Do not cache unresolved addresses for connections

- [HBASE-16194](#) - Should count in MSLAB chunk allocation into heap size change when adding duplicate cells
- [HBASE-16195](#) - Should not add chunk into chunkQueue if not using chunk pool in HeapMemStoreLAB
- [HBASE-16284](#) - Unauthorized client can shut down the cluster
- [HBASE-16317](#) - Revert all ESAPI changes
- [HBASE-16318](#) - Fail build while rendering velocity template if dependency license is not in whitelist
- [HBASE-16318](#) - Consistently use the correct name for "Apache License, Version 2.0"
- [HBASE-16321](#) - Ensure no findbugs-jsr305
- [HBASE-16340](#) - Exclude Xerces implementation jars from coming in transitively
- [HBASE-16360](#) - TableMapReduceUtil addHBaseDependencyJars has the wrong class name for PrefixTreeCodec
- [HIVE-9570](#) - Investigate test failure on union_view.q
- [HIVE-10007](#) - Support qualified table name in analyze table compute statistics for columns
- [HIVE-10728](#) - Deprecate unix_timestamp(void) and make it deterministic
- [HIVE-11901](#) - StorageBasedAuthorizationProvider requires write permission on table for SELECT statements
- [HIVE-12556](#) - Ctrl-C in Beeline does not kill Tez query on HS2
- [HIVE-13160](#) - HS2 unable to load UDFs on startup when HMS is not ready
- [HIVE-13620](#) - Merge llap branch work to master
- [HIVE-13645](#) - Beeline needs null-guard around hiveVars and hiveConfVars read
- [HIVE-14296](#) - Session count is not decremented when HS2 clients do not shutdown cleanly
- [HIVE-14436](#) - Hive 1.2.1/Hitting "ql.Driver: FAILED: IllegalArgumentException Error"
- [HIVE-14519](#) - Multi insert query bug
- [HIVE-14538](#) - Beeline throws exceptions with parsing Hive configuration when using !sh statement
- [HIVE-14715](#) - Hive throws NumberFormatException with query with Null value
- [HIVE-14743](#) - ArrayIndexOutOfBoundsException - HBASE-backed views query with JOINS
- [HUE-2689](#) - Sub-workflow submitted from coordinator gets parent workflow graph
- [HUE-4541](#) - Fixing Hue job browser - Kerberos mutual authentication error in Hue
- [HUE-4635](#) - Fix duration on jobs page for running jobs
- [HUE-4804](#) - Download function of HTML widget breaks the display
- [HUE-4808](#) - Do not show the edit link for sub-workflows when submitted outside Hue
- [HUE-4809](#) - Add truststore parameters only if SSL is turned on
- [HUE-4809](#) - Only add truststore paths when they actually exist
- [IMPALA-3081](#) - Increase memory limit for TestWideRow
- [IMPALA-3311](#) - Fix string data coming out of aggs in subplans
- [IMPALA-3575](#) - Add retry to back end connection request and rpc timeout
- [IMPALA-3678](#) - Fix migration of predicates into union operands with an order by + limit.
- [IMPALA-3682](#) - Do not retry unrecoverable socket creation errors
- [IMPALA-3687](#) - Fix test failure introduced by backporting
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3820](#) - Handle linkage errors while loading Java UDFs in Catalog
- [IMPALA-3930](#) - Fix shuffle insert hint with constant partition exprs
- [IMPALA-3940](#) - Fix getting column stats through views
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4049](#) - Fix empty batch handling NLJ build side
- [OOZIE-2068](#) - Configuration as part of sharelib
- [OOZIE-2347](#) - Remove unnecessary new Configuration()/new jobConf() calls from Oozie
- [OOZIE-2555](#) - Oozie SSL enable setup does not return port for admin -servers
- [OOZIE-2567](#) - HCat connection is not closed while getting hcat credentials
- [OOZIE-2589](#) - CompletedActionXCommand is hardcoded to wrong priority
- [OOZIE-2649](#) - Cannot override sub-workflow configuration property if defined in parent workflow XML
- [PIG-3807](#) - Pig creates wrong schema after dereferencing nested tuple fields with sorts
- [SPARK-8428](#) -Fix integer overflows in TimSort

- [SPARK-12339](#) - Added a null check that was removed in
- [SPARK-13242](#) - codegen fallback when there many branches

Issues Fixed in CDH 5.7.3

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.7.3:

- [FLUME-2821](#) - KafkaSourceUtil Can Log Passwords at Info remove logging of security related data in older releases.
- [FLUME-2913](#) - Don't strip SLF4J from imported classpaths.
- [FLUME-2922](#) - Sync SequenceFile.Writer before calling hflush
- [HADOOP-8751](#) - NPE in Token.toString() when Token is constructed using null identifier.
- [HADOOP-11361](#) - Fix a race condition in MetricsSourceAdapter.updateJmxCache.
- [HADOOP-11901](#) - BytesWritable fails to support 2G chunks due to integer overflow.
- [HADOOP-12659](#) - Incorrect usage of configuration parameters in token manager of KMS.
- [HADOOP-13263](#) - Reload cached groups in background after expiry.
- [HADOOP-13381](#) - KMS clients should use KMS Delegation Tokens from current UGI.
- [HADOOP-13457](#) - Remove hardcoded absolute path for shell executable.
- [HDFS-6434](#) - Default permission for creating file should be 644 for WebHdfs/HttpFS.
- [HDFS-7597](#) - DelegationTokenIdentifier should cache the TokenIdentifier to UGI mapping.
- [HDFS-8008](#) - Support client-side back off when the datanodes are congested.
- [HDFS-9276](#) - Failed to Update HDFS Delegation Token for long running application in HA mode.
- [HDFS-9466](#) - TestShortCircuitCache#testDataXceiverCleansUpSlotsOnFailure is flaky.
- [HDFS-9939](#) - Increase DecompressorStream skip buffer size.
- [HDFS-10512](#) - VolumeScanner may terminate due to NPE in DataNode.reportBadBlocks.
- [MAPREDUCE-6442](#) - Stack trace is missing when error occurs in client protocol provider's constructor.
- [MAPREDUCE-6473](#) - Job submission can take a long time during Cluster initialization.
- [MAPREDUCE-6675](#) - TestJobImpl.testUnusableNode failed
- [YARN-4459](#) - container-executor should only kill process groups.
- [YARN-4784](#) - Fairscheduler: defaultQueueSchedulingPolicy should not accept FIFO.
- [YARN-4866](#) - FairScheduler: AMs can consume all vcores leading to a livelock when using FAIR policy.
- [YARN-4878](#) - Expose scheduling policy and max running apps over JMX for Yarn queues.
- [YARN-5077](#) - Fix FSLeafQueue#getFairShare() for queues with zero fairshare.
- [YARN-5272](#) - Handle queue names consistently in FairScheduler.
- [HBASE-14963](#) - Remove use of Guava Stopwatch from HBase client code.
- [HBASE-15621](#) - Suppress Hbase SnapshotHFile cleaner error messages when a snapshot is going on.
- [HBASE-15808](#) - Reduce potential bulk load intermediate space usage and waste.
- [HBASE-16135](#) - PeerClusterZnode under rs of removed peer may never be deleted
- [HBASE-16207](#) - can't restore snapshot without "Admin" permission
- [HBASE-16227](#) - [Shell] Column value formatter not working in scans. Tested : manually using shell.
- [HBASE-16288](#) - HFile intermediate block level indexes might recurse forever creating multi TB files.
- [HBASE-16319](#) - Fix TestCacheOnWrite after HBASE-16288.
- [HIVE-11432](#) - Hive macro gives same result for different arguments.
- [HIVE-11487](#) - Add getNumPartitionsByFilter api in metastore api.
- [HIVE-11980](#) - Follow up on HIVE-11696, exception is thrown from CTAS from the table with table-level serde is Parquet while partition-level serde is JSON.
- [HIVE-12277](#) - Hive macro results on macro_duplicate.q different after adding ORDER BY.
- [HIVE-12635](#) - Hive should return the latest HBase cell timestamp as the row timestamp value.
- [HIVE-13043](#) - Reload function has no impact to function registry.
- [HIVE-13090](#) - Hive metastore crashes on NPE with ZooKeeperTokenStore.
- [HIVE-13372](#) - Hive Macro overwritten when multiple macros are used in one column.
- [HIVE-13704](#) - Do not call DistCp.execute() instead of DistCp.run().

- [HIVE-13749](#) - Memory leak in Hive Metastore.
- [HIVE-13884](#) - Disallow queries in HMS fetching more than a configured number of partitions
- [HIVE-14055](#) - directSql - getting the number of partitions is broken.
- [HIVE-14187](#) - JDOPersistenceManager objects remain cached if MetaStoreClient#close is not called.
- [HIVE-14209](#) - Add some logging info for session and operation management.
- [HIVE-14298](#) - NPE could be thrown in HMS when an ExpressionTree could not be made from a filter.
- [HIVE-14359](#) - Hive on Spark might fail in HS2 with LDAP authentication in a kerberized cluster.
- [HIVE-14457](#) - Partitions in encryption zone are still trashed though an exception is returned.
- [HUE-3481](#) - [assist] Do not sort the columns by name, instead use the creation order.
- [HUE-3842](#) - [core] HTTP 500 while emptying Hue 3.9 trash directory.
- [HUE-3845](#) - [sentry] Sometimes see group as editable on role section.
- [HUE-3880](#) - [core] Add importlib directly for Python 2.6.
- [HUE-3988](#) - [search] Support schemaless collections.
- [HUE-3999](#) - [oozie] list_oozie_workflow page should not break in case of bad json from oozie.
- [HUE-4265](#) - [beeswax] Bring back show preview in the assist.
- [HUE-4300](#) - [fb] Avoid double file listing call on folder search.
- [HUE-4333](#) - [core] Properly reset API_CACHE on failover.
- [HUE-4477](#) - [security] Select All is not filtering out the non visible roles from the selection .
- [HUE-4493](#) - [oozie] Fix sync-workflow action when Workflow includes sub-workflow.
- [HUE-4515](#) - [oozie] Remove oozie.bundle.application.path from properties when rerunning workflow.
- [IMPALA-3711](#) - Remove unnecessary privilege checks in getDbsMetadata().
- [IMPALA-3915](#) - Register privilege and audit requests when analyzing resolved table refs.
- [OOZIE-2391](#) - spark-opts value in workflow.xml is not parsed properly.
- [OOZIE-2537](#) - SqoopMain does not set up log4j properly.
- [SOLR-7280](#) - BackportLoad cores in sorted order and tweak coreLoadThread counts to improve cluster stability on restarts.
- [SOLR-9236](#) - AutoAddReplicas will append an extra /tlog to the update log location on replica failover.
- [SPARK-14963](#) - [YARN] Using recoveryPath if NM recovery is enabled.
- [SPARK-16505](#) - [YARN] Optionally propagate error during shuffle service startup.
- [SQOOP-2561](#) - Special Character removal from Column name as avro data results in duplicate column and fails the import.
- [SQOOP-2906](#) - Optimization of AvroUtil.toAvroIdentifier.
- [SQOOP-2971](#) - OraOop does not close connections properly.
- [SQOOP-2995](#) - Backward incompatibility introduced by Custom Tool options.

Issues Fixed in CDH 5.7.2

CDH 5.7.2 fixes the following issues.

Kerberized HS2 with LDAP authentication fails in a multi-domain LDAP case

In CDH 5.7, Hive introduced a feature to support HS2 with Kerberos plus LDAP authentication; but it broke compatibility with multi-domain LDAP cases on CDH 5.7.x and C5.8.x versions.

Affected Versions: CDH 5.7.1, CDH 5.8.0, and CDH 5.8.1

Fixed in Versions: CDH 5.7.2 and higher, CDH 5.8.2 and higher

Bug: [HIVE-13590](#).

Workaround: None.

Oozie

Oozie Web Console returns 500 error when Oozie server runs on JDK 8u75 or higher

Bug: [OOZIE-2533](#)

Cloudera Bug: CDH-40362

The Oozie Web Console returns a 500 error when the Oozie server is running on JDK 8u75 and higher. The Oozie server still functions, and you can use the Oozie command line, REST API, Java API, or the Hue Oozie Dashboard to review status of those jobs.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.7.2:

- [FLUME-1899](#) - Make SpoolDir work with subdirectories
- [FLUME-2910](#) - AsyncHBaseSink: Failure callbacks should log the exception that caused them
- [FLUME-2918](#) - Speed up TaildirSource on directories with many files
- [HADOOP-8934](#) - Shell command ls should include sort options
- [HADOOP-10971](#) - Add -C flag to make `hadoop fs -ls` print filenames only
- [HADOOP-11409](#) - FileContext.getFileContext can stack overflow if default fs misconfigured
- [HADOOP-11432](#) - Fix SymlinkBaseTest#testCreateLinkUsingPartQualPath2.
- [HADOOP-12787](#) - KMS SPNEGO sequence does not work with WebHDFS
- [HADOOP-12841](#) - Update s3-related properties in core-default.xml.
- [HADOOP-12901](#) - Add warning log when KMSClientProvider cannot create a connection to the KMS server.
- [HADOOP-12963](#) - Allow using path style addressing for accessing the S3 endpoint.
- [HADOOP-13079](#) - Add -q option to Ls to print ? instead of non-printable characters
- [HADOOP-13132](#) - Handle ClassCastException on AuthenticationException in LoadBalancingKMSClientProvider
- [HADOOP-13155](#) - Implement TokenRenewer to renew and cancel delegation tokens in KMS
- [HADOOP-13251](#) - Authenticate with Kerberos credentials when renewing KMS delegation token
- [HADOOP-13255](#) - KMSClientProvider should check and renew TGT when doing delegation token operations
- [HDFS-8581](#) - ContentSummary on / skips further counts on yielding lock
- [HDFS-8829](#) - Make SO_RCVBUF and SO_SNDBUF size configurable for DataTransferProtocol sockets and allow configuring auto-tuning
- [HDFS-9085](#) - Show renewer information in DelegationTokenIdentifier#toString
- [HDFS-9259](#) - Make SO_SNDBUF size configurable at DFSClient side for hdfs write scenario.
- [HDFS-9365](#) - Balancer does not work with the HDFS-6376 HA setup.
- [HDFS-9405](#) - Warmup NameNode EDEK caches in background thread
- [HDFS-9700](#) - DFSClient and DFSOutputStream should set TCP_NODELAY on sockets for DataTransferProtocol
- [HDFS-9732](#) - Improve DelegationTokenIdentifier.toString() for better logging
- [HDFS-9805](#) - Add server-side configuration for enabling TCP_NODELAY for DataTransferProtocol and default it to true
- [HDFS-10360](#) - DataNode may format directory and lose blocks if current/VERSION is missing
- [HDFS-10381](#) - DataStreamer DataNode exclusion log message should be warning
- [HDFS-10396](#) - Using -diff option with DistCp may get "Comparison method violates its general contract" exception
- [HDFS-10481](#) - HTTPFS server should correctly impersonate as end user to open file
- [HDFS-10516](#) - Fix bug when warming up EDEK cache of more than one encryption zone
- [HDFS-10525](#) - Fix NPE in CacheReplicationMonitor#rescanCachedBlockMap
- [MAPREDUCE-6558](#) - multibyte delimiters with compressed input files generate duplicate records
- [MAPREDUCE-6577](#) - MR AM unable to load native library without MR_AM_ADMIN_USER_ENV set
- [MAPREDUCE-6635](#) - Unsafe long to int conversion in UncompressedSplitLineReader and IndexOutOfBoundsExpection
- [MAPREDUCE-6701](#) - Application Master log unavailable when clicking JobHistory's AM logs link
- [YARN-2605](#) - [RM HA] Rest API endpoints doing redirect incorrectly.
- [YARN-4812](#) - TestFairScheduler#testContinuousScheduling fails intermittently.
- [YARN-4916](#) - Revert "TestNMProxy.testNMProxyRPCRetry fails
- [YARN-5048](#) - DelegationTokenRenewer#skipTokenRenewal may throw NPE
- [HBASE-11625](#) - Reading datablock throws "Invalid HFile block magic" and can not switch to hdfs checksum
- [HBASE-13532](#) - Make UnknownScannerException less scary by giving more information in the exception string.

- [HBASE-14644](#) - Region in transition metric is broken
- [HBASE-14818](#) - user_permission does not list namespace permissions
- [HBASE-15236](#) - Inconsistent cell reads over multiple bulk-loaded HFiles
- [HBASE-15439](#) - getMaximumAllowedTimeBetweenRuns in ScheduledChore ignores the TimeUnit
- [HBASE-15465](#) - userPermission returned by getUserPermission() for the selected namespace does not have namespace set
- [HBASE-15496](#) - Throw RowTooBigException only for user scan/get
- [HBASE-15698](#) - Increment TimeRange not serialized to server
- [HBASE-15746](#) - Remove extra RegionCoprocesor preClose() in RSRpcServices#closeRegion
- [HBASE-15791](#) - Improve javadoc around ScheduledChore
- [HBASE-15811](#) - Batch Get after batch Put does not fetch all cells
- [HBASE-15872](#) - Split TestWALProcedureStore
- [HBASE-15873](#) - ACL for snapshot restore / clone is not enforced
- [HBASE-15925](#) - provide default values for hadoop compat module related properties that match default hadoop profile.
- [HBASE-16034](#) - Fix ProcedureTestingUtility#LoadCounter.setMaxProclD()
- [HBASE-16056](#) - Procedure v2 - fix master crash for FileNotFound
- [HBASE-16093](#) - Fix splits failed before creating daughter regions leave meta inconsistent
- [HIVE-7443](#) - Fix HiveConnection to communicate with Kerberized Hive JDBC server and alternative JDKs
- [HIVE-9486](#) - Use session classloader instead of application loader
- [HIVE-9499](#) - hive.limit.query.max.table.partition makes queries fail on non-partitioned tables
- [HIVE-10685](#) - Alter table concatenate operator will cause duplicate data
- [HIVE-10925](#) - Non-static threadlocals in metastore code can potentially cause memory leak
- [HIVE-11031](#) - ORC concatenation of old files can fail while merging column statistics
- [HIVE-11243](#) - Changing log level in Utilities.getBaseWork
- [HIVE-11747](#) - Unnecessary error log is shown when executing a "INSERT OVERWRITE LOCAL DIRECTORY" cmd in the embedded mode
- [HIVE-11827](#) - STORED AS AVRO fails SELECT COUNT(*) when empty
- [HIVE-12742](#) - NULL table comparison within CASE does not work as previous hive versions
- [HIVE-12958](#) - Make embedded Jetty server more configurable
- [HIVE-13285](#) - ORC concatenation may drop old files from moving to final path
- [HIVE-13462](#) - HiveResultSetMetaData.getPrecision() fails for NULL columns
- [HIVE-13590](#) - Kerberized HS2 with LDAP auth enabled fails in multi-domain LDAP case
- [HIVE-13736](#) - View's input/output formats are TEXT by default.
- [HIVE-13932](#) - Hive SMB Map Join with small set of LIMIT failed with NPE
- [HIVE-13953](#) - Issues in HiveLockObject equals method
- [HIVE-13991](#) - Union All on view fail with no valid permission on underneath table
- [HIVE-14006](#) - Hive query with UNION ALL fails with ArrayIndexOutOfBoundsException.
- [HIVE-14015](#) - SMB MapJoin failed for Hive on Spark when kerberized
- [HIVE-14098](#) - Logging task properties and environment variables might contain passwords
- [HIVE-14118](#) - Make the alter partition exception more meaningful
- [HUE-2678](#) - [jobbrowser] Read Spark job data from Spark History Server API
- [HUE-3197](#) - [oozie] Decision node support in external Workflow graph
- [HUE-3520](#) - [jb] Use impersonation to access JHS if security is enabled
- [HUE-3521](#) - [core] Provide a force_username_uppercase option
- [HUE-3526](#) - [useradmin] Fix LDAP tests for force_username_uppercase
- [HUE-3688](#) - [oozie] Fix TestEditor.test_workflow_dependencies unit test
- [HUE-3700](#) - [core] Support force_username_lowercase and ignore_username_case for all Auth backends
- [HUE-3802](#) - [oozie] Fix HS2 action on SSL enabled cluster
- [HUE-3805](#) - [oozie] Add support for oozie schema 0.4 in dashboard graph for external workflows

- [HUE-3808](#) - [core] Offer to live turn on/off debug level
- [HUE-3821](#) - [pig] Logs are never returned on running script
- [HUE-3822](#) - [pig] Display logs when found
- [HUE-3861](#) - [core] Upgrade Django Axes to 1.5
- [HUE-3866](#) - [core] Hue CPU reaches ~100% usage while uploading files with SSL to HTTPFS/WebHDFS
- [HUE-3908](#) - [useradmin] Ignore (objectclass=*) filter when searching for LDAP users
- [HUE-3923](#) - [core] Simplify force debug logic option
- [HUE-4005](#) - [oozie] Remove oozie.coord.application.path from properties when rerunning workflow
- [HUE-4006](#) - [oozie] Create new deployment directory when coordinator or bundle is copied
- [HUE-4007](#) - [oozie] Fix deployment_dir for the bundle in oozie example fixtures
- [HUE-4021](#) - [libsolr] Allow customization of the Solr path in ZooKeeper
- [HUE-4023](#) - [useradmin] update AuthenticationForm to allow activated users to login
- [HUE-4061](#) - [job] Job attempt logs not appearing for running jobs
- [HUE-4087](#) - [jobbrowser] Unable to kill jobs with Resource Manager HA enabled
- [HUE-4092](#) - [security] Can't type any / in the HDFS ACLs path input
- [HUE-4113](#) - [Pig] Hue breaks when user has only access to pig app
- [HUE-4134](#) - [liboozie] Avoid logging truststore credentials
- [HUE-4202](#) - [job] Enable offset param for fetching jobbrowser logs
- [HUE-4215](#) - [yarn] Reset API_CACHE on logout
- [HUE-4227](#) - [yarn] Fix unittest for MR API Cache
- [HUE-4238](#) - [doc2] Ignore history docs in find_jobs_with_no_doc during sync documents
- [HUE-4252](#) - [core] Handle 307 redirect from YARN upon standby failover
- [HUE-4258](#) - [job] Close and pool Spark History Server connections
- [IMPALA-1928](#) - Fix Thrift client transport wrapping order
- [IMPALA-2660](#) - Respect auth_to_local configs from hdfs configs
- [IMPALA-3276](#) - Consistently handle pin failure in BTS::PrepareForRead()
- [IMPALA-3369](#) - Add ALTER TABLE SET COLUMN STATS statement.
- [IMPALA-3441](#) - Impala should not crash for invalid avro serialized data
- [IMPALA-3499](#) - Split catalog update
- [IMPALA-3502](#) - Fix race in the coordinator while updating filter routing table
- [IMPALA-3633](#) - Cancel fragment if coordinator is gone
- [IMPALA-3732](#) - Handle string length overflow in Avro files
- [IMPALA-3745](#) - Corrupt encoded values in parquet files can cause crashes
- [IMPALA-3751](#) - Fix clang build errors and warnings
- [IMPALA-3754](#) - Fix TestParquet.test_corrupt_rle_counts flakiness
- [OOZIE-2314](#) - Unable to kill old instance child job by workflow or coord rerun by Launcher
- [OOZIE-2329](#) - Make handling yarn restarts configurable
- [OOZIE-2330](#) - Spark action should take the global jobTracker and nameNode configs by default and allow file and archive elements
- [OOZIE-2345](#) - Parallel job submission for forked actions
- [OOZIE-2436](#) - Fork/join workflow fails with oozie.action.yarn.tag must not be null
- [OOZIE-2481](#) - Add YARN_CONF_DIR in the Shell action
- [OOZIE-2504](#) - Create a log4j.properties under HADOOP_CONF_DIR in Shell Action
- [OOZIE-2511](#) - SubWorkflow missing variable set from option if config-default is present in parent workflow
- [OOZIE-2533](#) - Oozie Web UI gives Error 500 with Java 8u91
- [SENTRY-1175](#) - Improve usability of URI privileges when granting URIs
- [SENTRY-1201](#) - Sentry ignores database prefix for MSCK statement
- [SENTRY-1252](#) - grantServerPrivilege and revokeServerPrivilege should treat "*" and "ALL" as synonyms when action is not explicitly specified
- [SENTRY-1265](#) - Sentry service should not require a TGT as it is not talking to other kerberos services as a client

- [SENTRY-1292](#) - Reorder DBModelAction EnumSet
- [SENTRY-1293](#) - Avoid converting string permission to Privilege object
- [SENTRY-1311](#) - Improve usability of URI privileges by supporting mixed use of URIs with and without scheme
- [SENTRY-1320](#) - truncate table db_name.table_name fails
- [SOLR-7178](#) - OverseerAutoReplicaFailoverThread compares Integer objects using ==
- [SOLR-8451](#) - We should not call method.abort in HttpSolrClient and HttpSolrCall#remoteQuery should not close streams
- [SOLR-8497](#) - Merge indexes should mark its directories as done rather than keep them around in the directory cache.
- [SOLR-8691](#) - Cache index fingerprints per searcher
- [SOLR-9053](#) - Upgrade commons-fileupload to 1.3.1, fixing a potential vulnerability
- [SPARK-13278](#) - [CORE] Launcher fails to start with JDK 9 EA
- [SPARK-14391](#) - [LAUNCHER] Fix launcher communication test
- [SPARK-15067](#) - [YARN] YARN executors are launched with fixed perm gen size
- [SPARK-15165](#) - [SPARK-15205] [SQL] Introduce place holder for comments in generated code
- [SQOOP-2846](#) - Sqoop Export with update-key failing for avro data file
- [SQOOP-2864](#) - ClassWriter chokes on column names containing double quotes
- [SQOOP-2920](#) - Sqoop performance deteriorates significantly on wide datasets; sqoop 100% on CPU

Issues Fixed in CDH 5.7.1

CDH 5.7.1 fixes the following issues.

Apache HBase

The `ReplicationCleaner` process can abort if its connection to ZooKeeper is inconsistent

Bug: [HBASE-15234](#)

If the connection with ZooKeeper is inconsistent, the `ReplicationCleaner` may abort, and the following event is logged by the HMaster:

```
WARN org.apache.hadoop.hbase.replication.master.ReplicationLogCleaner: Aborting
ReplicationLogCleaner
because Failed to get list of replicators
```

Unprocessed WALs accumulate.

The `seekBefore()` method calculates the size of the previous data block by assuming that data blocks are contiguous, and HFile v2 and higher store Bloom blocks and leaf-level INode blocks with the data. As a result, reverse scans do not work when Bloom blocks or leaf-level INode blocks are present when HFile v2 or higher is used.

Workaround: Restart the HMaster occasionally. The `ReplicationCleaner` restarts if necessary and process the unprocessed WALs.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.7.1:

- [AVRO-1781](#) - Schema.parse is not thread safe
- [FLUME-2781](#) - Kafka Channel with parseAsFlumeEvent=true should write data as is, not as flume events
- [FLUME-2891](#) - Revert FLUME-2712 and FLUME-2886
- [FLUME-2897](#) - AsyncHBase sink NPE when Channel.getTransaction() fails
- [HADOOP-7139](#) - Allow appending to existing SequenceFiles
- [HADOOP-7817](#) - RawLocalFileSystem.append() should give FSDataOutputStream with accurate .getPos()
- [HADOOP-11321](#) - copyToLocal cannot save a file to an SMB share unless the user has Full Control permissions
- [HADOOP-11687](#) - Ignore x-* and response headers when copying an Amazon S3 object
- [HADOOP-12668](#) - Support excluding weak Ciphers in HttpServer2 through ssl-server.conf
- [HADOOP-12825](#) - Log slow name resolutions

- [HADOOP-12954](#) - Add a way to change `hadoop.security.token.service.use_ip`
- [HADOOP-12972](#) - `Lz4Compressor#getLibraryName` returns the wrong version number
- [HDFS-3519](#) - Checkpoint upload may interfere with a concurrent `saveNamespace`
- [HDFS-6520](#) - `hdfs fsck` passes invalid length value when creating `BlockReader`
- [HDFS-7600](#) - Refine `hdfs` admin classes to reuse common code
- [HDFS-8142](#) - `DistributedFileSystem` encryption zone commands should resolve relative paths
- [HDFS-8211](#) - `DataNode` UUID is always null in the JMX counter.
- [HDFS-8496](#) - Calling `stopWriter()` with `FSDatasetImpl` lock held may block other threads
- [HDFS-8855](#) - `Webhdfs` client leaks active `NameNode` connections
- [HDFS-9549](#) - `TestCacheDirectives#testExceedsCapacity` is flaky
- [HDFS-9589](#) - Block files which have been hardlinked should be duplicated before the `DataNode` appends to the them
- [HDFS-9949](#) - Add a test case to ensure that the `DataNode` does not regenerate its UUID when a storage directory is cleared
- [HDFS-10223](#) - `peerFromSocketAndKey` performs SASL exchange before setting connection timeouts
- [HDFS-10267](#) - Extra "synchronized" on `FsDatasetImpl#recoverAppend` and `FsDatasetImpl#recoverClose`
- [HDFS-10324](#) - Trash directory in an encryption zone should be pre-created with correct permissions
- [HDFS-10344](#) - `DistributedFileSystem#getTrashRoots` should skip encryption zone that does not have `.Trash`
- [MAPREDUCE-4785](#) - `TestMRApp` occasionally fails
- [MAPREDUCE-6297](#) - Task Id of the failed task in diagnostics should link to the task page
- [MAPREDUCE-6333](#) - `TestEvents`, `TestAMWebServicesTasks`, `TestAppController` are broken due to [MAPREDUCE-6297](#)
- [MAPREDUCE-6384](#) - Add the last reporting reducer info for too many fetch failure diagnostics
- [MAPREDUCE-6388](#) - Remove deprecation warnings from `JobHistoryServer` classes
- [MAPREDUCE-6485](#) - Create a new task attempt with failed map task priority if in-progress attempts are unassigned
- [MAPREDUCE-6513](#) - MR job got hanged forever when one NM unstable for some time
- [MAPREDUCE-6535](#) - `TaskID` default constructor results in NPE on `toString()`
- [MAPREDUCE-6580](#) - Test failure: `TestMRJobsWithProfiler`
- [YARN-2871](#) - `TestRMRestart#testRMRestartGetApplicationList` sometime fails in trunk
- [YARN-3104](#) - Fixed RM to not generate new AMRM tokens on every heartbeat between rolling and activation
- [YARN-3493](#) - RM fails to come up with error "Failed to load/recover state" when mem settings are changed
- [YARN-3695](#) - `ServerProxy` (`NMProxy`, etc.) shouldn't retry forever for non network exception
- [YARN-4168](#) - Fixed a failing test `TestLogAggregationService.testLocalFileDeletionOnDiskFull`
- [YARN-4414](#) - `Nodemanager` connection errors are retried at multiple levels
- [YARN-4579](#) - Allow `DefaultContainerExecutor` container log directory permissions to be configurable
- [YARN-4629](#) - Distributed shell breaks under strong security
- [YARN-4639](#) - Remove dead code in `TestDelegationTokenRenewer` added in [YARN-3055](#)
- [YARN-4717](#) - `TestResourceLocalizationService.testPublicResourceInitializesLocalDir` fails Intermittently due to `IllegalArgumentException` from cleanup
- [YARN-4795](#) - `ContainerMetrics` drops records
- [YARN-4916](#) - `TestNMProxy.testNMProxyRPCRetry` fails
- [HBASE-15234](#) - Don't abort `ReplicationLogCleaner` on ZooKeeper errors
- [HBASE-15271](#) - Spark bulk load should write to temporary location and then rename on success
- [HBASE-15349](#) - Update surefire version to 2.19.1
- [HBASE-15405](#) - Fix PE logging and wrong defaults in help message
- [HBASE-15456](#) - `CreateTableProcedure/ModifyTableProcedure` needs to fail when there is no family in table descriptor
- [HBASE-15479](#) - No more garbage or beware of autoboxing
- [HBASE-15582](#) - `SnapshotManifestV1` too verbose when there are no regions
- [HBASE-15591](#) - `ServerCrashProcedure` not yielding
- [HBASE-15592](#) - Print Procedure WAL content

- [HBASE-15622](#) - Superusers does not consider the keytab credentials
- [HBASE-15622](#) - Superusers does not consider the keytab credentials
- [HBASE-15673](#) - Fix latency metrics for multiGet. - Also fixes some stuff in help text
- [HBASE-15707](#) - ImportTSV bulk output does not support tags with hfile.format.version=3
- [HIVE-6099](#) - Multi insert does not work properly with distinct count
- [HIVE-10303](#) - HIVE-9471 broke forward compatibility of ORC files
- [HIVE-10313](#) - Literal Decimal ExprNodeConstantDesc should contain value of HiveDecimal instead of String
- [HIVE-10396](#) - moredecimal_precision2.q test is failing on trunk
- [HIVE-10636](#) - CASE comparison operator rotation optimization
- [HIVE-11054](#) - Handle varchar/char partition columns in vectorization
- [HIVE-11097](#) - HiveInputFormat uses String.startsWith to compare splitPath and PathToAliases
- [HIVE-11369](#) - Mapjoins in HiveServer2 fail when jmxremote is used
- [HIVE-11408](#) - HiveServer2 is leaking ClassLoaders when add jar / temporary functions are used due to constructor caching in Hadoop ReflectionUtils
- [HIVE-11427](#) - Location of temporary table for CREATE TABLE SELECT broken by HIVE-7079
- [HIVE-11590](#) - AvroDeserializer is very chatty
- [HIVE-11919](#) - Hive Union Type Mismatch
- [HIVE-12481](#) - Occasionally "Request is a replay" will be thrown from HS2
- [HIVE-12506](#) - SHOW CREATE TABLE command creates a table that does not work for RCFile format
- [HIVE-12568](#) - Provide an option to specify network interface used by Spark remote client [Spark Branch]
- [HIVE-12616](#) - NullPointerException when spark session is reused to run a mapjoin
- [HIVE-12706](#) - Incorrect output from from_utc_timestamp()/to_utc_timestamp when local timezone has DST
- [HIVE-12941](#) - Unexpected result when using MIN() on struct with NULL in first field
- [HIVE-13082](#) - Enable constant propagation optimization in query with left semi join
- [HIVE-13082](#) - Enable constant propagation optimization in query with left semi join
- [HIVE-13115](#) - MetaStore Direct SQL getPartitions call fail when the columns schemas for a partition are null
- [HIVE-13200](#) - Aggregation functions returning empty rows on partitioned columns
- [HIVE-13217](#) - Replication for HoS mapjoin small file needs to respect dfs.replication.max
- [HIVE-13243](#) - Hive drop table on encryption zone fails for external tables
- [HIVE-13251](#) - hive can't read the decimal in AVRO file generated from previous version
- [HIVE-13286](#) - Query ID is being reused across queries
- [HIVE-13295](#) - Improvement to LDAP search queries in HS2 LDAP Authenticator
- [HIVE-13300](#) - Hive on spark throws exception for multi-insert with join
- [HIVE-13302](#) - direct SQL: cast to date doesn't work on Oracle
- [HIVE-13376](#) - HoS emits too many logs with application state
- [HIVE-13401](#) - Kerberized HS2 with LDAP auth enabled fails kerberos/delegation token authentication
- [HIVE-13410](#) - PerfLog metrics scopes not closed if there are exceptions on HS2
- [HIVE-13500](#) - Fix OOM with explan output being logged
- [HIVE-13527](#) - Using deprecated APIs in HBase client causes zookeeper connection leaks
- [HIVE-13530](#) - Hive on Spark throws Kryo exception in some cases
- [HIVE-13570](#) - Some queries with Union all fail when CBO is off
- [HIVE-13585](#) - Add counter metric for direct sql failures
- [HIVE-13632](#) - CDH39911Hive failing on insert empty array into parquet table
- [HIVE-13657](#) - Spark driver stderr logs should appear in hive client logs
- [HUE-3171](#) - Fix vertical resize handle for queries with long descriptions
- [HUE-3171](#) - Long descriptions doesn't wrap and headers from table follows with horizontal scroll
- [HUE-3221](#) - Styling on column stats popup leaks on the tables page
- [HUE-3293](#) - Prevent document matching query error when going one home 1
- [HUE-3293](#) - Fix mis-switching to new home page when new editor is on
- [HUE-3293](#) - Move new editor flag to desktop

- [HUE-3303](#) - PostgreSQL requires data update and alter table operations in separate transactions
- [HUE-3310](#) - Prevent browsing job designs by API
- [HUE-3334](#) - Update test, now se send empty query instead of error
- [HUE-3334](#) - Skip checking for multi queries if there is no semi colon
- [HUE-3350](#) - Reverse browsing link to use the correct version of the editor
- [HUE-3398](#) - Filter out sessions with empty guid or secret key
- [HUE-3434](#) - Logs of finished Oozie workflow are not displayed
- [HUE-3436](#) - Retain old dependencies when saving a workflow
- [HUE-3437](#) - PamBackend does not honor ignore_username_case
- [HUE-3459](#) - Put stat popover on top
- [HUE-3459](#) - Fixed Flexbox for IE10
- [HUE-3459](#) - Use fixed positioning for assist panel
- [HUE-3459](#) - Revert sticky assist
- [HUE-3459](#) - Fix issue with single panel in metastore and new editor
- [HUE-3459](#) - Clear the height interval on update
- [HUE-3459](#) - Assist doesn't stretch to the end of the page in the old editors
- [HUE-3471](#) - Set the assist database on design update
- [HUE-3471](#) - Assist does not show the DB from the saved query
- [HUE-3476](#) - Clear any running intervals after closing the stats popover
- [HUE-3480](#) - Impala refresh pop-over won't close after assist action while open
- [HUE-3506](#) - Limit length of comments on table page
- [HUE-3511](#) - Reduce flickering of action icons when moving the pointer across several entries
- [HUE-3523](#) - Modify find_jobs_with_no_doc method to exclude jobs with no name
- [HUE-3528](#) - Call correct metrics api to avoid 500 error
- [HUE-3543](#) - Timeout prevents refreshing of the Assist tables/dbs
- [HUE-3594](#) - Smarter DOM based XSS filter on hashes
- [HUE-3601](#) - Spinner not positioned correctly
- [HUE-3613](#) - Empty div elements are added when scrolling the DB assist panel
- [HUE-3614](#) - Scrolling on assist in old editor also scrolls the editor
- [HUE-3637](#) - Avoid decode errors on attribute values
- [HUE-3650](#) - Notify of caught errors in the watch logs process
- [HUE-3651](#) - Upgrade Moment.js
- [HUE-3704](#) - Force enable notebook permissions
- [HUE-3716](#) - Add gen-py paths to hue.pth
- [HUE-3725](#) - 'SparkJob' object has no attribute 'amHostHttpAddress'
- [HUE-3731](#) - Send database on Impala refresh with invalidate
- [HUE-3741](#) - Display field validation errors on create table wizard
- [HUE-3800](#) - Job attempt logs not appearing for some Oozie jobs
- [HUE-3819](#) - Make the upload and create icons not disappear under 1180px
- [IMPALA-2076](#) - Correct execution time tracking for DataStreamSender
- [IMPALA-2502](#) - Don't redundantly repartition grouping aggregations
- [IMPALA-2892](#) - Buffered-tuple-stream-ir.cc is not cross-compiled
- [IMPALA-3133](#) - Wrong privileges after a REVOKE ALL ON SERVER statement
- [IMPALA-3139](#) - Fix drop table statement to not drop views and vice versa
- [IMPALA-3141](#) - Send dummy filters when filter production is disabled
- [IMPALA-3194](#) - Allow queries materializing scalar type columns in RC/sequence files
- [IMPALA-3220](#) - Skip logging empty ScannerContext's stream in parse error
- [IMPALA-3236](#) - Increase timeout for runtime filter tests
- [IMPALA-3238](#) - Avoid log spam for very large hash tables
- [IMPALA-3245](#), [IMPALA-3305](#): Fix crash with global filters when NUM_NODES=1

- [IMPALA-3269](#) - Remove authz checks on default table location in CTAS queries
- [IMPALA-3285](#) - Fix ASAN failure in webservice-test
- [IMPALA-3317](#) - Fix crash in sorter when spilling zero-length strings
- [IMPALA-3334](#) - Fix some bugs in query options' parsing.
- [IMPALA-3367](#) - Ensure runtime filters tests run on 3 nodes
- [IMPALA-3378](#), [IMPALA-3379](#): fix various JNI issues
- [IMPALA-3385](#) - Fix crashes on accessing error_log
- [IMPALA-3395](#) - Old HT filter code uses wrong expr type
- [IMPALA-3396](#) - Fix ConcurrentTimerCounter unit test "TimerCounterTest" failure.
- [IMPALA-3412](#) - Fix CHAR codegen crash in tuple comparator
- [IMPALA-3420](#) - Set IMPALA_THRIFT_VERSION patch level to +4
- [KITE-1108](#) - Add optional retry feature to loadSolr morphline command
- [KITE-1114](#) - Kite CLI json-import HDFS temp file path not multiuser safe
- [OOZIE-2429](#) - TestEventGeneration test is flakey
- [OOZIE-2466](#) - Repeated failure of TestMetricsInstrumentation.testSamplers
- [OOZIE-2486](#) - TestSLAEventsGetForFilterJPAExecutor is flakey
- [OOZIE-2490](#) - Oozie can't set hadoop.security.token.service.use_ip
- [SENTRY-922](#) - INSERT OVERWRITE DIRECTORY permission not working correctly
- [SENTRY-1112](#) - Change default value of "sentry.hive.server" to empty string
- [SENTRY-1164](#) - testCaseSensitivity test failure on a real cluster and also a minor improvements to testConcurrentClients to run locally. (Anne Yu, reviewed by Haohao).
- [SENTRY-1169](#) - MetastorePlugin#renameAuthzObject log message prints oldpathname as newpathname
- [SENTRY-1184](#) - Clean up HMSPaths.renameAuthzObject
- [SENTRY-1190](#) - IMPORT TABLE silently fails if Sentry is enabled
- [SOLR-6631](#) - DistributedQueue spinning on calling zookeeper getChildren()
- [SOLR-6879](#) - Have an option to disable autoAddReplicas temporarily for all collections.
- [SOLR-7493](#) - Requests aren't distributed evenly if the collection isn't present locally. Merges r1683946 and r1683948 from trunk.
- [SOLR-8551](#) - Make collection deletion more robust.
- [SOLR-8683](#) - Tune down stream closed logging
- [SOLR-8720](#) - ZkController#publishAndWaitForDownStates should use #publishNodeAsDown.
- [SOLR-8771](#) - Multi-threaded core shutdown creates executor per core
- [SOLR-8855](#) - The HDFS BlockDirectory should not clean up it's cache on shutdown.
- [SOLR-8856](#) - Do not cache merge or 'read once' contexts in the hdfs block cache.
- [SOLR-8857](#) - HdfsUpdateLog does not use configured or new default number of version buckets and is hard coded to 256.
- [SOLR-8869](#) - Optionally disable printing field cache entries in SolrFieldCacheMBean
- [SPARK-4452](#) - Shuffle data structures can starve others on the same thread for memory
- [SPARK-12614](#) - Don't throw non fatal exception from ask
- [SPARK-13622](#) - Issue creating level db for YARN shuffle service
- [SPARK-14242](#) - Avoid copy in compositeBuffer for frame decoder
- [SPARK-14290](#) - Avoid significant memory copy in Netty's tran...
- [SPARK-14363](#) - Fix executor OOM due to memory leak in the Sorter
- [SPARK-14477](#) - Allow custom mirrors for downloading artifacts in build/mvn
- [SPARK-14679](#) - Fix UI DAG visualization OOM.
- [SQOOP-2847](#) - Sqoop --incremental + missing parent --target-dir reports success with no data

Issues Fixed in CDH 5.7.0

CDH 5.7.0 fixes the following issues.

Apache Flume

TailDirSource throws FileNotFoundException Exception if ~/.flume directory is not created already

Bug: [FLUME-2773](#)

This fix ensures that any missing parent directories in the `positionFile` path (either default or user given input) are always created.

flume_env script should handle JVM parameters like -javaagent -agentpath -agentlib

Bug: [FLUME-2763](#)

This fix enables the `flume_env` script to handle JVM parameters such as `-javaagent`, `-agentpath` and `-agentlib`.

Kafka channel timeout property is overridden by default value

Bug: [FLUME-2734](#)

When the Kafka channel timeout property is passed to the Kafka consumer internally, it does not work as expected. It is overridden by the default value or the value specified by the `.timeout` property, which is undocumented. Now the `kafka.consumer.timeout.ms` value specified in the configuration takes effect like other Kafka consumer properties.

Apache Hadoop

ReplicationMonitor can infinitely loop in BlockPlacementPolicyDefault#chooseRandom()

Bug: [HDFS-4937](#)

Cloudera Bug: CDH-34043

When a large number of nodes are removed by refreshing node lists, the network topology is updated and the replication monitor thread may get stuck in the `while` loop of `chooseRandom()`.

Clean up temporary files after fsimage transfer failures

Bug: [HDFS-7373](#)

Cloudera Bug: CDH-19177

When an `fsimage` (or checkpoint) transfer fails, a temporary file is left in each storage directory. If the namespace is large, these files can take up a large amount of space.

Lease recovery should return true if the lease can be released and the file can be closed

Bug: [HDFS-8576](#)

Cloudera Bug: CDH-37212

`FSNamesystem#recoverLease` should return `true` when a lease is recovered both explicitly and implicitly—that is, when a lease recovery is successful and the file is closed, and also when a file is closed and the lease is released *without* a recovery.

fsck does not list correct file path when bad replicas or blocks are in a snapshot

Bug: [HDFS-9231](#)

Cloudera Bug: CDH-32221

When blocks are corrupt in a snapshot, the `fsck` command lists the original directory and not the snapshot directory. This happens even when the original file is deleted. The specific commands are `fsck -list-corruptfileblocks` and `fsck -list-corruptfileblocks -includeSnapshots`.

Make DataStreamer#block thread safe and verify generationStamp in commitBlock

Bug: [HDFS-9289](#)

Cloudera Bug: CDH-33723

When the client calls `updatePipeline`, a block might commit with an old `generationStamp`, causing replicas to look corrupt.

*Delayed heartbeat processing causes storm of subsequent heartbeats***Bug:** [HDFS-9305](#)**Cloudera Bug:** CDH-33589

The NameNode usually handles DataNode heartbeats quickly, but can be delayed for various reasons, such as a long garbage collection or lock contention. After the NameNode recovers, the DataNode sends a storm of heartbeat messages in a tight loop which, in a big cluster, can overload the NameNode and make cluster recovery difficult.

*FSImage may get corrupted after deleting snapshot***Bug:** [HDFS-9406](#)**Cloudera Bug:** CDH-33224

When deleting a snapshot that contains the last record of a given INode, the fsimage may become corrupt because the create list of the snapshot diff in the previous snapshot and the child list of the parent INodeDirectory are not cleaned.

Apache HBase

See also [Known Issues In CDH 5.7.0](#) on page 1973.

*Potential data loss after a RegionServerAbortedException***Bug:** [HBASE-13895](#)

If the master attempts to assign a region while handling a RegionServer abort, the returned RegionServerAbortedException is handled as though the region had been cleanly taken offline, so the new assignment is allowed to proceed. If the region is opened in its new location before WAL replay has completed, the replayed edits are ignored, or are later played back on top of new edits that happened after the region was opened. In either case, data can be lost.

Workaround: None.*Data loss can occur if a table has more than 2,147,483,647 columns***Bug:** [HBASE-15133](#)

Data loss can occur if a table has more than 2,147,483,647 (`Integer.MAX_INT`) columns, because some key variable types are INT rather than LONG.

Workaround: Adjust your schema to use fewer than `Integer.MAX_INT` columns.*Delete operations that occur during a region merge may be eclipsed by new Put operations***Bug:** [HBASE-13938](#)

The master's timestamp is not used when sending `hbase:meta` edits on region merges, so correct ordering of new region additions and old region deletes is not assured and data loss can occur if edits are applied in the wrong order.

Workaround: None.*RPC handler / task monitoring seems to be broken after 0.98***Bug:** [HBASE-14674](#)

After pluggable RPC scheduler, the way the tasks work for the handlers got changed. We no longer list idle RPC handlers in the tasks, but we register them dynamically to TaskMonitor through CallRunner. However, the IPC readers are still registered the old way (meaning that idle readers are listed as tasks, but not idle handlers).

From the javadoc of MonitoredRPCHandlerImpl, it seems that we are NOT optimizing the allocation for the MonitoredTask anymore, but instead allocate one for every RPC call breaking the pattern

*Conflicts between HBase Balancer and hbase:meta reassignment***Bug:** [HBASE-14536](#)

If `hbase:meta` is assigned to a `RegionServer` that becomes unavailable, and the HBase balancer has scheduled but not completed a plan to move `hbase:meta` to a different `RegionServer`, the `hbase:meta` becomes unassigned.

Workaround: None.

Regions can fail to transition in a write-heavy cluster with a small number of read handlers

Bug: [HBASE-13635](#)

On a write-heavy cluster configured with a small number of read handlers, all requests that are not mutations are sent to the read handlers, including `ReportRegionInTransition` requests. If these requests time out, the `RegionServer` is assumed to be unavailable, and the regions cannot transition correctly.

Workaround: None.

In a secured environment, when a RegionServer is stopped, znodes may not be cleaned up correctly

Bug: [HBASE-14581](#)

When a `RegionServer` process is stopped, the `zkcli` command is invoked to delete its `znodes`. In a secure cluster, the `zkcli` command does not authenticate to ZooKeeper and the deletion fails. This problem occurs because the `REGIONSERVER_OPTS` environment variable is not correctly passed when invoking the `zkcli` command.

Workaround: None.

Delays in RegionServer responses can cause a region to be closed indefinitely

Bug: [HBASE-14407](#)

Handling of region assignment by the master has a flaw when `RegionServer` responses are delayed due to network delays, system load, or other reasons. This flaw can cause the master to close a region indefinitely.

Workaround: Restart the `RegionServer` to force the region to be reassigned.

When a RegionServer crashes, replication peers can crash due to inode exhaustion from old WALs

Bug: [HBASE-14621](#)

The fix for [HBASE-12865](#) ensures that `loadWALsFromQueues` attempts a retry when the replication source version is changed while loading the replication queue. However, the fix introduced a bug in `ReplicationLogCleaner` that causes an infinite loop when a `RegionServer` crashes. As a result, old WALs are not cleaned up. In a cluster under high load, the inode limit on the replication peer `RegionServer` can be exhausted, causing the `RegionServer` to crash.

Workaround: None.

When a RegionServer crashes, cell-level visibility tags may be lost during WAL replay

Bug: [HBASE-15218](#)

When reading cells after a `RegionServer` crash, the `KeyValueCodec` and the `WallCellCodec` both use `NoTagsKeyValue`, which does not preserve visibility tags.

Workaround: None.

Column is not deleted if you do not pass the visibility label

Bug: [HBASE-14761](#)

If a column was created or modified with a visibility label, and you attempt to delete it without passing the visibility label, the column is not deleted. It is not visible using a `Scan` operation, but is visible using a raw `Scan`, and is marked with `deleteColumn`.

Workaround: None.

If multiple users are configured with the role `hbase.superuser`, an attempt to connect to a secure ZooKeeper instance fails

Bug: [HBASE-14425](#)

The `hbase.superuser` configuration option is a comma-separated list of users. A bug in the code to connect to a secure ZooKeeper causes the list to be evaluated as a single value, so a list of multiple users fails because no username matches the comma-separated list.

Workaround: Only specify a single user in the `hbase.superuser` configuration option.

Region split request audits are performed against the `hbase` user instead of the requesting user

Bug: [HBASE-14475](#)

When checking whether the requesting user has permission to perform a region split, the `hbase` user's permissions are checked instead of those of the requesting user. Due to this bug, `CREATE` is sufficient for the split, rather than `CREATE` and `ADMIN`. Because `CREATE` permissions are also sufficient for flushes and compactions, this issue is not severe in most environments.

Workaround: None.

Incorrect timestamp checking causes unpredictable deletes with `VisibilityScanDeleteChecker`.

Bug: [HBASE-13635](#)

Incorrect timestamp checking when `VisibilityScanDeleteChecker` is used causes unpredictable results when deleting cells. In some cases, the timestamp is deleted but the cell contents are not deleted. In other cases, a request to delete an entire row or to delete a version results in only a single cell being deleted.

Workaround: None.

A `BulkLoad` of an HFile with tags that requires splits does not preserve the tags

Bug: [HBASE-15035](#)

When an HFile is created with cell tags and is imported into HBase using a bulk load, the tags are present as expected when the HFile is loaded into a single region. However, if the bulk load spans multiple regions, the original HFile is automatically split into a set of HFiles corresponding to each of the regions the original HFile covers. Tags, including ACLs, TTLs, and MOB pointers, are not copied to the split files.

Workaround: None.

Restoring a snapshot from a table in a user-defined namespace causes a `URISyntaxException`

Bug: [HBASE-14578](#)

A table in a user-defined namespace uses a colon between the namespace and the table name (for instance, `example_ns:users`). This colon is interpreted incorrectly when restoring from a snapshot.

Workaround: None.

The `list_snapshots` HBase shell command shows all snapshots, regardless of the user's permission to view them

Bug: [HBASE-12552](#)

A user with no privileges to interact with a snapshot can list the snapshot using the `list_snapshots` HBase shell command.

Workaround: None.

`ExportSnapshot` or `DistCp` operations may fail on the Amazon `s3a://` protocol

Bug: None.

`ExportSnapshot` or `DistCP` operations may fail on AWS when using certain JDK 8 versions, due to an incompatibility between AWS Java SDK 1.9.x and the `joda-time` date-parsing module.

Workaround: Use `joda-time` 2.8.1 or higher, which is included in AWS Java SDK 1.10.1 or higher.

If HDFS is restarted while HBase is running, WALs being replicated may not close correctly

Bug: [HBASE-15019](#)

The RegionServer receiving the replicated WALs has no mechanism to be notified to perform a recovery if HDFS is restarted on the source cluster.

Workaround: Restart the RegionServer to force the master to trigger the lease recovery during WAL splitting.

The permissions of the .top/ directory are not explicitly set during LoadIncrementalHFiles operations

Bug: [HBASE-14005](#)

Permissions are not explicitly set on the .top/ directory created during LoadIncrementalHFiles. The permissions should be set the same as the .bottom/ and .tmp/ directories.

Workaround: None.

Nonfatal errors in the FSHLog subsystem are incorrectly logged as fatal errors

Bug: [HBASE-14042](#)

If an IOException causes a log roll to be requested, it is logged as a fatal event, although it should be logged as a warning.

Workaround: None.

FuzzyRowFilter may omit some rows if multiple fuzzy keys are present

Bug: [HBASE-14269](#)

If you use the FuzzyRowFilter for Scan operations, and you have multiple fuzzy keys, some rows may be omitted from the RowTracker.

Workaround: None.

The prefix-tree module is not automatically included in MapReduce jobs

Bug: [HBASE-15152](#)

JARs for prefix-tree module are not automatically included in YarnChildren processes. This causes a ClassNotFoundException.

Workaround: Manually add the prefix-tree JARs to the classpath if needed.

Values of some metrics may appear to be negative

Bug: [HBASE-12961](#)

Some metric value are stored in integers, and cannot accommodate real-world values. This causes metric values to appear to be negative.

Workaround: None.

The HBase Shell cannot handle Scan filters which contain non-UTF8 characters

Bug: [HBASE-15032](#)

The HBase Shell incorrectly handles filter strings which contain non-UTF8 characters.

Workaround: None.

Reverse scans do not work when Bloom blocks or leaf-level inode blocks are present

Bug: [HBASE-14283](#)

Because the seekBefore() method calculates the size of the previous data block by assuming that data blocks are contiguous, and HFile v2 and higher store Bloom blocks and leaf-level inode blocks with the data, reverse scans do not work when Bloom blocks or leaf-level inode blocks are present when HFile v2 or higher is used.

Workaround: None.

Apache Hive

Fix regression in bind and search logic for Hive external LDAP authentication

Bug: [HIVE-12885](#)

Cloudera Bug: CDH-35075

Fixes a regression in LDAP bind and search authentication from CDH 5.5.0.

Some queries using LEFT SEMI JOIN fail with IndexOutOfBoundsException

Bug: [HIVE-13082](#)

Cloudera Bug: CDH-37515

Some queries using LEFT SEMI JOIN fail with IndexOutOfBoundsException. Constant propagation optimization for these queries is now enabled.

BETWEEN predicate is not functioning correctly with predicate pushdown on Parquet tables

Bug: [HIVE-13039](#)

Cloudera Bug: CDH-37322

BETWEEN becomes exclusive in Parquet table when predicate pushdown (PPD) is enabled for Parquet tables, leading to potential incorrect results.

Performance degradation when running Hive queries against wide tables with Sentry enabled

Bug: [SENTRY-1007](#)

Cloudera Bug: CDH-35908

Fixes a performance regression due to inefficient authorization checks in the Sentry Hive binding for Hive tables that are wide (more than 100 columns).

Optionally cancel queries after configurable timeout waiting on compilation lock

Bug: [HIVE-12431](#)

Cloudera Bug: CDH-34693

Adds a new configuration option, `hive.server2.compile.lock.timeout`, that cancels queries if they are waiting for the compile lock for more than this amount of time. This applies only to queries waiting on compilation and does not cancel queries that are being compiled. By default, the timeout is unlimited.

Hue

The Hive Sample Table, customer, Cannot be Queried

Bug: [HUE-3040](#)

Cloudera Bug: CDH-35054

Users can now query the customer table in the Hive editor.

Apache Impala

For the list of Impala fixed issues, see [Issues Fixed in Impala for CDH 5.7.0](#) on page 2422.

See also [Apache Impala Known Issues](#) on page 1995 for issues that are known but not resolved yet.

MapReduce

MapReduce Rolling Upgrades To and From CDH 5.6.0 Fail

Cloudera Bug: CDH-38587

Users can now safely use rolling upgrade from releases CDH 5.6.0 and lower to CDH 5.7.0.

Cloudera Search

Reordered updates cause leaders and replicas to become out of sync

Bug: [SOLR-8586](#), [SOLR-8691](#)

Solr relied on checking leader/replica document synchronization by comparing the last 100 updates on the leader and replica for significant overlap, and then applying any missing updates from the leader. In certain cases, document updates could be significantly reordered, resulting in mismatches in the index, even when the last 100 documents

matched. Solr now implements hashing over the versions of all the documents to check for synchronization, eliminating a class of errors in which replicas and leaders could become out of sync.

Apache Sentry

Fixed Sentry Oracle upgrade script

Bug: [SENTRY-1066](#)

This fixes previous Sentry upgrade issues with Oracle (ORA-0955).

Tables with non-HDFS locations break Hive Metastore startup

Bug: [SENTRY-1044](#)

Tables with non-HDFS locations cause the HDFS/Sentry plugin to enter an invalid state and fail with the exception, Could not create Initial AuthzPaths or HMSHandler !!.

URI check is now case-sensitive

Bug: [SENTRY-968](#)

Sentry no longer ignores case when validating privileges for URIs.

TRUNCATE on empty partitioned table in Hive fails

Bug: [SENTRY-826](#)

PathsUpdate.parsePath(path) will throw an NPE when parsing relative paths

Bug: [SENTRY-1002](#)

Sentry now skips relative paths (that is, paths without a fully qualified scheme) rather than failing with a NPE.

The Sentry Server should not be updated if the CREATE/DROP operations fail

Bug: [SENTRY-1008](#)

Previously, even if a CREATE TABLE operation failed, the Sentry Server would still be updated with a path to the table. This has been fixed.

SimpleDBProviderBackend should retry the authorization process

Bug: [SENTRY-902](#)

This fix includes corrections to the retry logic to remove recursive calls and include a wait time between retries when authorization fails.

Support Hive RELOAD by updating the classpath for Sentry

Bug: [SENTRY-1003](#)

When Hive issues the RELOAD command, Sentry now gets the updated auxiliary JAR path from the `hive.reloadable.aux.jars.path` property.

RealTimeGet with explicit ids can bypass document-level authorization

Bug: [SENTRY-989](#)

Users can no longer bypass security by guessing the document IDs for the `RealTimeGet` command.

Updated Apache Shiro dependency

Bug: [SENTRY-1054](#)

External partitions referenced by more than one table can cause some unexpected behavior with Sentry HDFS sync

Bug: [SENTRY-953](#)

INSERT INTO command no longer requires URI privilege on partition location under table

Bug: [SENTRY-1095](#)

The checks on the Hive INSERT INTO command have been relaxed. The INSERT INTO command adds location information to the partition description. Although this requires a check on URI privileges, in this case location information can be generated even if the partition is under the table directory.

Improvement to the SentryAuthFilter error message when authentication fails

Bug: [SENTRY-1060](#)

Avoid logging all DataNucleus queries when debug logging is enabled

Bug: [SENTRY-945](#)

Logging DataNucleus queries when debugging can fill up 2 GB of log file space in less than five minutes.

getGroup and getUser should always return original HDFS values for paths that are not managed by Sentry

Bug: [SENTRY-936](#)

Paths that do not correspond to Hive metastore objects should not be affected by HDFS/Sentry sync.

Exceptions in MetastoreCacheInitializer should not prevent HMS from starting up

Bug: [SENTRY-957](#)

Instead of only throwing a runtime exception, this fix ensures failed tasks are first retried.

Set maximum message size for Thrift messages

Bug: [SENTRY-904](#)

This ensures that security scans and unformatted messages do not bring down the Sentry server by going out of bounds.

Allow SentryAuthorization setter path always fall through and update HDFS

Bug: [SENTRY-988](#)

Setting HDFS rules on Sentry-managed HDFS paths should not affect original HDFS rules

Bug: [SENTRY-944](#)

Removing and modifying ACLs on Sentry-managed paths that correspond to Hive metastore objects should not affect the original HDFS access rules.

Fix inconsistency in column-level privileges

Bug: [SENTRY-847](#)

If you have column-level privileges for a table, the SHOW columns operation should not require extra table-level privileges.

Performance Improvements

Improved performance for filtering Hive SHOW commands

Bug: [SENTRY-565](#)

HiveAuthzBinding has been improved to reduce the number of RPC calls when filtering SHOW commands.

Improved Sentry column-level performance for wide tables

Bug: [SENTRY-1007](#)

Apache Spark

Spark SQL does not support the char type

Spark SQL does not support the `char` type (fixed-length strings). Like unions, tables with such fields cannot be created from or read by Spark.

Spark SQL statements that can result in table partition metadata changes may fail

Because Spark does not have access to Sentry data, it may not know that a user has permissions to execute an operation and instead fail it. SQL statements that can result in table partition metadata changes, for example, "ALTER TABLE" or "INSERT", may fail.

Cloudera Bug: CDH-33446

Certain Spark MLlib features not supported

The following Spark MLlib features are now supported:

- `spark.ml`
- ML pipeline APIs

Streaming incompatibility between Spark 1.2 and 1.3

Applications built as a JAR with dependencies ("uber JAR") must be built for the specific version of Spark running on the cluster.

Cloudera Bug: CDH-26527

Workaround: Rebuild the JAR with the Spark dependencies in `pom.xml` pointing to the specific version of Spark running on the target cluster.

Spark SQL cannot retrieve data from a partitioned Hive table

When reading from a partitioned Hive table, Spark SQL cannot identify the column delimiter used and reads the full record as the first column entry.

Cloudera Bug: CDH-37189

Workaround: Contact Cloudera Support for information on how to deploy a patch to resolve the issue.

Tasks that fail due to YARN preemption can cause job failure

Bug: [SPARK-8167](#)

Tasks that are running on preempted executors will count as FAILED with an `ExecutorLostFailure`.

Apache Sqoop

Oracle connector not working with lowercase columns

Bug: [SQOOP-2723](#)

The Oracle connector now works with lowercase columns.

Run only one map task attempt during export

Bug: [SQOOP-2712](#)

In most scenarios, running multiple map task attempts by default when performing an export is not required. The default is now one map task attempt during export operations.

Do not dump data on error in TextExportMapper by default

Bug: [SQOOP-2651](#)

Dumping data in the `TextExportMapper` might unintentionally leak sensitive information to logs. The `enableDataDumpOnError` key is set to `false` by default. A user can set the value to `true` to intentionally write the data to the log.

Support of glob paths during export

Bug: [SQOOP-1281](#)

Glob paths are now supported for export.

Sqoop should support importing from table with column names containing some special characters

Bug: [SQOOP-2387](#)

Sqoop supports some special characters in column names. The specific list of characters depends on those supported for a particular database.

Avro export ignores --columns option

Bug: [SQOOP-1369](#)

AvroExportMapper now supports the `--columns` option to restrict the columns to export.

JDK

Java 8 (updates 60 and higher) has problems with joda-time and S3 requests

Bug: [SPARK-11413](#)

Cloudera Bug: CDH-31245

Versions of Java 1.8, from update 60 and higher (jdk1.8.0_60++), cause S3 to fail because joda-time cannot format time zones.

Issues Fixed in CDH 5.6.x

The following topics describe issues fixed in CDH 5.6.x, from newest to oldest release. You can also review [What's New In CDH 5.6.x](#) on page 1834 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.6.1

CDH 5.6.1 fixes the following issues.

Apache Hadoop

FSImage may get corrupted after deleting snapshot

Bug: [HDFS-9406](#)

Cloudera Bug: CDH-33224

When deleting a snapshot that contains the last record of a given INode, the fsmimage may become corrupt because the create list of the snapshot diff in the previous snapshot and the child list of the parent `INodeDirectory` are not cleaned.

Apache HBase

The ReplicationCleaner process can abort if its connection to ZooKeeper is inconsistent

Bug: [HBASE-15234](#)

If the connection with ZooKeeper is inconsistent, the `ReplicationCleaner` may abort, and the following event is logged by the HMaster:

```
WARN org.apache.hadoop.hbase.replication.master.ReplicationLogCleaner: Aborting
ReplicationLogCleaner
because Failed to get list of replicators
```

Unprocessed WALs accumulate.

The `seekBefore()` method calculates the size of the previous data block by assuming that data blocks are contiguous, and HFile v2 and higher store Bloom blocks and leaf-level INode blocks with the data. As a result, reverse scans do not work when Bloom blocks or leaf-level INode blocks are present when HFile v2 or higher is used.

Workaround: Restart the HMaster occasionally. The `ReplicationCleaner` restarts if necessary and process the unprocessed WALs.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.6.1:

- [FLUME-2632](#) - High CPU on KafkaSink
- [FLUME-2712](#) - Optional channel errors slows down the Source to Main channel event rate
- [FLUME-2781](#) - Kafka Channel with `parseAsFlumeEvent=true` should write data as is, not as flume events

- [FLUME-2886](#) - Optional Channels can cause OOMs
- [FLUME-2891](#) - Revert FLUME-2712 and FLUME-2886
- [FLUME-2897](#) - AsyncHBase sink NPE when Channel.getTransaction() fails
- [HADOOP-7139](#) - Allow appending to existing SequenceFiles
- [HADOOP-7817](#) - RawLocalFileSystem.append() should give FSDataOutputStream with accurate .getPos()
- [HADOOP-11171](#) - Enable using a proxy server to connect to S3a
- [HADOOP-11321](#) - copyToLocal cannot save a file to an SMB share unless the user has Full Control permissions
- [HADOOP-11687](#) - Ignore x-* and response headers when copying an Amazon S3 object
- [HADOOP-11722](#) - Some Instances of Services using ZKDelegationTokenSecretManager go down when old token cannot be deleted
- [HADOOP-12240](#) - Fix tests requiring native library to be skipped in non-native profile
- [HADOOP-12280](#) - Skip unit tests based on maven profile rather than NativeCodeLoader.isNativeCodeLoaded
- [HADOOP-12604](#) - Exception may be swallowed in KMSClientProvider
- [HADOOP-12605](#) - Fix intermittent failure of TestIPC.testIpcWithReaderQueuing
- [HADOOP-12668](#) - Support excluding weak Ciphers in HttpServer2 through ssl-server.conf
- [HADOOP-12699](#) - TestKMS#testKMSProvider intermittently fails during 'test rollover draining'
- [HADOOP-12715](#) - TestValueQueue#testgetAtMostPolicyALL fails intermittently
- [HADOOP-12718](#) - Incorrect error message by fs -put local dir without permission
- [HADOOP-12736](#) - TestTimedOutTestsListener#testThreadDumpAndDeadlocks sometimes times out
- [HADOOP-12825](#) - Log slow name resolutions
- [HADOOP-12954](#) - Add a way to change hadoop.security.token.service.use_ip
- [HADOOP-12972](#) - Lz4Compressor#getLibraryName returns the wrong version number
- [HDFS-6520](#) - hdfs fsck passes invalid length value when creating BlockReader
- [HDFS-8211](#) - DataNode UUID is always null in the JMX counter
- [HDFS-8496](#) - Calling stopWriter() with FSDataSetImpl lock held may block other threads
- [HDFS-8576](#) - Lease recovery should return true if the lease can be released and the file can be closed
- [HDFS-8785](#) - TestDistributedFileSystem is failing in trunk
- [HDFS-8855](#) - Webhdfs client leaks active NameNode connections
- [HDFS-9264](#) - Minor cleanup of operations on FsVolumeList#volumes
- [HDFS-9289](#) - Make DataStreamer#block thread safe and verify genStamp in commitBlock
- [HDFS-9347](#) - Invariant assumption in TestQuorumJournalManager.shutdown() is wrong
- [HDFS-9350](#) - Avoid creating temporary strings in Block.toString() and getBlockName()
- [HDFS-9358](#) - TestNodeCount#testNodeCount timed out
- [HDFS-9514](#) - TestDistributedFileSystem.testDFSClientPeerWriteTimeout failing; exception being swallowed
- [HDFS-9576](#) - HTrace: collect position/length information on read operations
- [HDFS-9589](#) - Block files which have been hardlinked should be duplicated before the DataNode appends to the them
- [HDFS-9612](#) - DistCp worker threads are not terminated after jobs are done
- [HDFS-9655](#) - NN should start JVM pause monitor before loading fsimage.
- [HDFS-9688](#) - Test the effect of nested encryption zones in HDFS downgrade
- [HDFS-9701](#) - DN may deadlock when hot-swapping under load
- [HDFS-9721](#) - Allow Delimited PB OIV tool to run upon fsimage that contains INodeReference
- [HDFS-9949](#) - Add a test case to ensure that the DataNode does not regenerate its UUID when a storage directory is cleared
- [HDFS-10223](#) - peerFromSocketAndKey performs SASL exchange before setting connection timeouts
- [HDFS-10267](#) - Extra "synchronized" on FsDataSetImpl#recoverAppend and FsDataSetImpl#recoverClose
- [MAPREDUCE-4785](#) - TestMRApp occasionally fails
- [MAPREDUCE-6460](#) - TestRMContainerAllocator.testAttemptNotFoundCausesRMCommunicatorException fails
- [MAPREDUCE-6528](#) - Memory leak for HistoryFileManager.getJobSummary()
- [MAPREDUCE-6580](#) - Test failure: TestMRJobsWithProfiler

- [MAPREDUCE-6620](#) - Jobs that did not start are shown as starting in 1969 in the JHS web UI
- [YARN-2749](#) - Fix some testcases from TestLogAggregationService fails in trunk
- [YARN-2871](#) - TestRMRestart#testRMRestartGetApplicationList sometime fails in trunk
- [YARN-2902](#) - Killing a container that is localizing can orphan resources in the DOWNLOADING state
- [YARN-3104](#) - Fixed RM to not generate new AMRM tokens on every heartbeat between rolling and activation
- [YARN-3446](#) - FairScheduler headroom calculation should exclude nodes in the blacklist
- [YARN-3727](#) - For better error recovery, check if the directory exists before using it for localization
- [YARN-4155](#) - TestLogAggregationService.testLogAggregationServiceWithInterval failing
- [YARN-4168](#) - Fixed a failing test TestLogAggregationService.testLocalFileDeletionOnDiskFull
- [YARN-4354](#) - Public resource localization fails with NPE
- [YARN-4380](#) - TestResourceLocalizationService.testDownloadingResourcesOnContainerKill fails intermittently
- [YARN-4393](#) - Fix intermittent test failure for TestResourceLocalizationService#testFailedDirsResourceRelease
- [YARN-4546](#) - ResourceManager crash due to scheduling opportunity overflow
- [YARN-4573](#) - Fix test failure in TestRMAppTransitions#testAppRunningKill and testAppKilledKilled
- [YARN-4613](#) - Fix test failure in TestClientRMService#testGetClusterNodes
- [YARN-4704](#) - TestResourceManager#testResourceAllocation() fails when using FairScheduler
- [YARN-4717](#) - TestResourceLocalizationService.testPublicResourceInitializesLocalDir fails Intermittently due to IllegalArgumentExpection from cleanup
- [HBASE-6617](#) - ReplicationSourceManager should be able to track multiple WAL paths
- [HBASE-14374](#) - Stuck FSHLog' issue to 1.1 Also includes HBASE-14807 TestWALLockup is flakey
- [HBASE-14759](#) - Avoid using Math.abs when selecting SyncRunner in FSHLog
- [HBASE-15019](#) - Replication stuck when HDFS is restarted
- [HBASE-15052](#) - Use EnvironmentEdgeManager in ReplicationSource
- [HBASE-15152](#) - Automatically include prefix-tree module in MR jobs if present
- [HBASE-15157](#) - Add *PerformanceTest for Append, CheckAnd* Reason: Bug Author: Stack Ref:
- [HBASE-15206](#) - Fix flakey testSplitDaughtersNotInMeta
- [HBASE-15213](#) - Fix increment performance regression caused by HBASE-8763 on branch-1.0
- [HBASE-15234](#) - Don't abort ReplicationLogCleaner on ZooKeeper errors
- [HBASE-15456](#) - CreateTableProcedure/ModifyTableProcedure needs to fail when there is no family in table descriptor
- [HBASE-15479](#) - No more garbage or beware of autoboxing
- [HBASE-15582](#) - SnapshotManifestV1 too verbose when there are no regions
- [HIVE-6099](#) - Multi insert does not work properly with distinct count
- [HIVE-7653](#) - Hive AvroSerDe does not support circular references in Schema
- [HIVE-9617](#) - UDF from_utc_timestamp throws NPE if the second argument is null
- [HIVE-10115](#) - HS2 running on a Kerberized cluster should offer Kerberos(GSSAPI) and Delegation token(DIGEST) when alternate authentication is enabled
- [HIVE-10213](#) - MapReduce jobs using dynamic-partitioning fail on commit
- [HIVE-10303](#) - HIVE-9471 broke forward compatibility of ORC files
- [HIVE-11054](#) - Handle varchar/char partition columns in vectorization
- [HIVE-11097](#) - HiveInputFormat uses String.startsWith to compare splitPath and PathToAliases
- [HIVE-11135](#) - Fix the Beeline set and save command in order to avoid the NullPointerException
- [HIVE-11285](#) - ObjectInspector for partition columns in FetchOperator in SMBJoin causes exception
- [HIVE-11288](#) - Avro SerDe InstanceCache returns incorrect schema
- [HIVE-11408](#) - HiveServer2 is leaking ClassLoaders when add jar / temporary functions are used due to constructor caching in Hadoop ReflectionUtils
- [HIVE-11427](#) - Location of temporary table for CREATE TABLE SELECT broken by HIVE-7079
- [HIVE-11488](#) - Combine the following jiras for "Support sessionId and queryId logging"Add sessionId and queryId info to HS2 log
- [HIVE-12456](#): QueryId can't be stored in the configuration of the SessionState since multiple queries can run in a single session

- [HIVE-11583](#) - When PTF is used over a large partitions result could be corrupted
- [HIVE-11590](#) - AvroDeserializer is very chatty
- [HIVE-11828](#) - beeline -f fails on scripts with tabs between column type and comment
- [HIVE-11919](#) - Hive Union Type Mismatch
- [HIVE-12315](#) - Fix Vectorized double divide by zero
- [HIVE-12354](#) - MapJoin with double keys is slow on MR
- [HIVE-12431](#) - Support timeout for compile lock
- [HIVE-12469](#) - Apache Commons Collections
- [HIVE-12506](#) - SHOW CREATE TABLE command creates a table that does not work for RCFile format
- [HIVE-12706](#) - Incorrect output from from_utc_timestamp()/to_utc_timestamp when local timezone has DST
- [HIVE-12790](#) - Metastore connection leaks in HiveServer2
- [HIVE-12885](#) - LDAP Authenticator improvements
- [HIVE-12941](#) - Unexpected result when using MIN() on struct with NULL in first field
- [HIVE-12946](#) - alter table should also add default scheme and authority for the location similar to create table
- [HIVE-13039](#) - BETWEEN predicate is not functioning correctly with predicate pushdown on Parquet table
- [HIVE-13055](#) - Add unit tests for HIVE-11512
- [HIVE-13065](#) - Hive throws NPE when writing map type data to a HBase backed table
- [HIVE-13082](#) - Enable constant propagation optimization in query with left semi join
- [HIVE-13200](#) - Aggregation functions returning empty rows on partitioned columns
- [HIVE-13243](#) - Hive drop table on encryption zone fails for external tables
- [HIVE-13251](#) - Hive can't read the decimal in AVRO file generated from previous version
- [HIVE-13286](#) - Query ID is being reused across queries
- [HIVE-13295](#) - Improvement to LDAP search queries in HS2 LDAP Authenticator
- [HIVE-13401](#) - Kerberized HS2 with LDAP auth enabled fails kerberos/delegation token authentication
- [HIVE-13527](#) - Using deprecated APIs in HBase client causes zookeeper connection leaks
- [HIVE-13570](#) - Some queries with Union all fail when CBO is off
- [HUE-3106](#) - [filebrowser] Add support for full paths in zip file uploads
- [HUE-3110](#) - [oozie] Fix bundle submission when coordinator points to multiple bundles
- [HUE-3132](#) - [core] Fix Sync Ldap users and groups for anonymous binds
- [HUE-3180](#) - [useradmin] Override duplicate username validation message
- [HUE-3185](#) - [oozie] Avoid extra API calls for parent information in workflow dashboard
- [HUE-3303](#) - [core] PostgreSQL requires data update and alter table operations in separate transactions
- [HUE-3310](#) - [jobsub] Prevent browsing job designs by API
- [HUE-3334](#) - [editor] Update test, now send empty query instead of error, skip checking for multi queries if there is no semicolon
- [HUE-3398](#) - [beeswax] Filter out sessions with empty guid or secret key
- [HUE-3436](#) - [oozie] Retain old dependencies when saving a workflow
- [HUE-3437](#) - [core] PamBackend does not honor ignore_username_case
- [HUE-3523](#) - [oozie] Modify find_jobs_with_no_doc method to exclude jobs with no name
- [HUE-3528](#) - [oozie] Call correct metrics api to avoid 500 error
- [HUE-3594](#) - [fb] Smarter DOM based XSS filter on hashes
- [HUE-3637](#) - [sqoop] Avoid decode errors on attribute values
- [HUE-3650](#) - [beeswax] Notify of caught errors in the watch logs process
- [HUE-3651](#) - [core] Upgrade Moment.js
- [IMPALA-852](#), [IMPALA-2215](#) - Analyze HAVING clause before aggregation
- [IMPALA-1092](#) - Fix estimates for trivial coord-only queries
- [IMPALA-1170](#) - Fix URL parsing when path contains '@'
- [IMPALA-1934](#) - Allow shell to retrieve LDAP password from shell cmd
- [IMPALA-2093](#) - Disallow NOT IN aggregate subqueries with a constant lhs expr
- [IMPALA-2184](#) - don't inline timestamp methods with try/catch blocks in IR

- [IMPALA-2425](#) - Broadcast join hint not enforced when low memory limit is set
- [IMPALA-2503](#) - Add missing String.format() arg in error message
- [IMPALA-2539](#) - Unmark collections slots of empty union operands
- [IMPALA-2554](#) - Change default buffer size for RPC servers and clients
- [IMPALA-2565](#) - Planner tests are flaky due to file size mismatches
- [IMPALA-2592](#) - DataStreamSender::Channel::CloseInternal() does not close the channel on an error
- [IMPALA-2599](#) - Pseudo-random sleep before acquiring kerberos ticket possibly not really pseudo-random
- [IMPALA-2711](#) - Fix memory leak in Rand()
- [IMPALA-2732](#) - Timestamp formats with non-padded values
- [IMPALA-2734](#) - Correlated EXISTS subqueries with HAVING clause return wrong results
- [IMPALA-2742](#) - Avoid unbounded MemPool growth with AcquireData()
- [IMPALA-2749](#) - Fix decimal multiplication overflow
- [IMPALA-2765](#) - Preserve return type of subexpressions substituted in isTrueWithNullSlots()
- [IMPALA-2788](#) - conv(bigint num, int from_base, int to_base) returns wrong result
- [IMPALA-2798](#) - Bring in AVRO-1617 fix and add test case for it
- [IMPALA-2818](#) - Fix cancellation crashes/hangs due to BlockOnWait() race
- [IMPALA-2820](#) - Support unquoted keywords as struct-field names
- [IMPALA-2832](#) - Fix cloning of FunctionCallExpr
- [IMPALA-2844](#) - Allow count(*) on RC files with complex types
- [IMPALA-2870](#) - Fix failing metadata.test_ddl.TestDdlStatements.test_create_table test
- [IMPALA-2894](#) - Move regression test into a different .test file
- [IMPALA-2906](#) - Fix an edge case with materializing TupleIsNotNullPredicates in analytic sorts
- [IMPALA-2914](#) - Fix DCHECK Check failed: HasDateOrTime()
- [IMPALA-2926](#) - Fix off-by-one bug in SelectNode::CopyRows()
- [IMPALA-2940](#) - Fix leak of dictionaries in Parquet scanner
- [IMPALA-3000](#) - Fix BitReader::Reset()
- [IMPALA-3034](#) - Verify all consumed memory of a MemTracker is always released at destruction time
- [IMPALA-3047](#) - Separate create table test with nested types
- [IMPALA-3054](#) - Disable probe side filters when spilling
- [IMPALA-3071](#) - Fix assignment of On-clause predicates belonging to an inner join
- [IMPALA-3085](#) - Unregister data sinks' MemTrackers at their Close() functions
- [IMPALA-3093](#) - ReopenClient() could NULL out 'client_key' causing a crash
- [IMPALA-3095](#) - Add configurable whitelist of authorized internal principals
- [IMPALA-3151](#) - Impala crash for avro table when casting to char data type
- [IMPALA-3194](#) - Allow queries materializing scalar type columns in RC/sequence files
- [KITE-1114](#) - Kite CLI json-import HDFS temp file path not multiuser safe, fix missing license header
- [OOZIE-2419](#) - HBase credentials are not correctly proxied
- [OOZIE-2466](#) - Repeated failure of TestMetricsInstrumentation.testSamplers
- [OOZIE-2486](#) - TestSLAEventsGetForFilterJPAExecutor is flakey
- [OOZIE-2490](#) - Oozie can't set hadoop.security.token.service.use_ip
- [SENTRY-748](#) - Improve test coverage of Sentry + Hive using complex views
- [SENTRY-835](#) - Drop table leaves a connection open when using metastorelistener
- [SENTRY-922](#) - INSERT OVERWRITE DIRECTORY permission not working correctly
- [SENTRY-972](#) - Include sentry-tests-hive hadoop test script in maven project
- [SENTRY-991](#) - Roles of Sentry Permission needs to be case insensitive
- [SENTRY-994](#) - SentryAuthorizationInfoX should override isSentryManaged
- [SENTRY-1002](#) - PathsUpdate.parsePath(path) will throw an NPE when parsing relative paths
- [SENTRY-1003](#) - Support "reload" by updating the classpath of Sentry function aux jar path during runtime
- [SENTRY-1007](#) - Sentry column-level performance for wide tables
- [SENTRY-1008](#) - Path should be not be updated if the create/drop table/partition event fails

- [SENTRY-1015](#) - Improve Sentry + Hive error message when user has insufficient privileges
- [SENTRY-1044](#) - Tables with non-hdfs locations breaks HMS startup
- [SENTRY-1169](#) - MetastorePlugin#renameAuthzObject log message prints oldpathname as newpathname
- [SENTRY-1184](#) - Clean up HMSPaths.renameAuthzObject
- [SOLR-6631](#) - DistributedQueue spinning on calling zookeeper getChildren()
- [SOLR-6820](#) - The sync on the VersionInfo bucket in DistributedUpdateProcessor#addDocument appears to be a large bottleneck when using replication
- [SOLR-7281](#) - Add an overseer action to publish an entire node as 'down'
- [SOLR-7332](#) - Seed version buckets with max version from index
- [SOLR-7493](#) - Requests aren't distributed evenly if the collection isn't present locally. Merges r1683946 and r1683948 from trunk
- [SOLR-7587](#) - TestSpellCheckResponse stalled and never timed out -- possible VersionBucket bug?
- [SOLR-7625](#) - Version bucket seed not updated after new index is installed on a replica
- [SOLR-8215](#) - Only active replicas should handle incoming requests against a collection
- [SOLR-8371](#) - Try and prevent too many recovery requests from stacking up and clean up some faulty cancel recovery logic
- [SOLR-8451](#) - We should not call method.abort in HttpSolrClient or HttpSolrCall#remoteQuery and HttpSolrCall#remoteQuery should not close streams
- [SOLR-8453](#) - Solr should attempt to consume the request inputStream on errors as we cannot count on the container to do it
- [SOLR-8575](#) - Fix HDFSLogReader replay status numbers and a performance bug where we can reopen FSDataInputStream too often
- [SOLR-8578](#) - Successful or not, requests are not always fully consumed by Solrj clients and we count on HttpClient or the JVM
- [SOLR-8615](#) - Just like creating cores, we should use multiple threads when closing cores
- [SOLR-8633](#) - DistributedUpdateProcess processCommit/deleteByQuery calls finish on DUP and SolrCmdDistributor, which violates the lifecycle and can cause bugs
- [SOLR-8683](#) - Always consume the full request on the server, not just in the case of an error, tune down stream closed logging
- [SOLR-8720](#) - ZkController#publishAndWaitForDownStates should use #publishNodeAsDown
- [SOLR-8771](#) - Multi-threaded core shutdown creates executor per core
- [SOLR-8855](#) - The HDFS BlockDirectory should not clean up it's cache on shutdown
- [SOLR-8856](#) - Do not cache merge or 'read once' contexts in the hdfs block cache
- [SOLR-8857](#) - HdfsUpdateLog does not use configured or new default number of version buckets and is hard coded to 256
- [SOLR-8869](#) - Optionally disable printing field cache entries in SolrFieldCacheMBean
- [SPARK-10859](#) - Predicates pushed to InmemoryColumnarTableScan are not evaluated correctly
- [SPARK-10914](#) - UnsafeRow serialization breaks when two machines have different Oups size
- [SPARK-11009](#) - RowNumber in HiveContext returns negative values in cluster mode
- [SPARK-11442](#) - Reduce numSlices for local metrics test of SparkListenerSuite
- [SPARK-12617](#) - Socket descriptor leak killing streaming app
- [SPARK-14477](#) - Allow custom mirrors for downloading artifacts in build/mvn
- [SQOOP-2847](#) - Sqoop --incremental + missing parent --target-dir reports success with no data

Issues Fixed in CDH 5.6.0

Apache Avro

Concurrent schema parsing can result in processes hanging because of an unsafe shared cache

Bug: [AVRO-1781](#)

Cloudera Bug: CDH-36057

Severity: High

Workaround: Use a global lock around all calls to `Schema.parse` and `Schema.Parser#parse` to ensure only one thread is parsing at a time.

Apache HBase

Values of some metrics may appear to be negative.

Bug: [HBASE-12961](#)

Some metric value are stored in integers, and cannot accommodate real-world values. This causes metric values to appear to be negative.

Workaround: None.

The HBase Shell cannot handle Scan filters which contain non-UTF8 characters.

Bug: [HBASE-15032](#)

The HBase Shell incorrectly handles filter strings which contain non-UTF8 characters.

Workaround: None.

HDFS

Checkpointing can fail due to an InvalidSignatureException in a secure cluster

Bug: [HDFS-7798](#)

Cloudera Bug: CDH-33411

Severity: High

Workaround: This problem occurs occasionally due to race condition. The error is transient, and a subsequent checkpoint may still succeed.

Issues Fixed in CDH 5.5.x

The following topics describe issues fixed in CDH 5.5.x, from newest to oldest release. You can also review [What's New In CDH 5.5.x](#) on page 1834 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.5.6

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.5.6:

- [FLUME-2797](#) - Use SourceCounter for SyslogTcpSource
- [FLUME-2844](#) - SpillableMemoryChannel must start ChannelCounter
- [HADOOP-10300](#) - Allow deferred sending of call responses
- [HADOOP-11031](#) - Design Document for Credential Provider API
- [HADOOP-12453](#) - Support decoding KMS Delegation Token with its own Identifier
- [HADOOP-12483](#) - Maintain wrapped SASL ordering for postponed IPC responses
- [HADOOP-12537](#) - S3A to support Amazon STS temporary credentials
- [HADOOP-12548](#) - Read s3a credentials from a Credential Provider
- [HADOOP-12609](#) - Fix intermittent failure of TestDecayRpcScheduler.
- [HADOOP-12723](#) - S3A: Add ability to plug in any AWSCredentialsProvider
- [HADOOP-12749](#) - Create a threadpoolexecutor that overrides afterExecute to log uncaught exceptions/errors
- [HADOOP-13034](#) - Log message about input options in distcp lacks some items
- [HADOOP-13317](#) - Add logs to KMS server-side to improve supportability
- [HADOOP-13353](#) - LdapGroupsMapping getPassword should not return null when IOException throws
- [HADOOP-13487](#) - Hadoop KMS should load old delegation tokens from ZooKeeper on startup
- [HADOOP-13526](#) - Add detailed logging in KMS for the authentication failure of proxy user
- [HADOOP-13558](#) - UserGroupInformation created from a Subject incorrectly tries to renew the Kerberos ticket
- [HADOOP-13638](#) - KMS should set UGI's Configuration object properly
- [HADOOP-13669](#) - KMS Server should log exceptions before throwing

- [HADOOP-13693](#) - Remove the message about HTTP OPTIONS in SPNEGO initialization message from kms audit log.
- [HDFS-4176](#) - EditLogTailer should call rollEdits with a timeout.
- [HDFS-6962](#) - ACLs inheritance conflict with umaskmode
- [HDFS-7210](#) - Avoid two separate RPC's namenode.append() and namenode.getFileInfo() for an append call from DFSClient
- [HDFS-7413](#) - Some unit tests should use NameNodeProtocols instead of FSNameSystem
- [HDFS-7415](#) - Move FSNameSystem.resolvePath() to FSDirectory
- [HDFS-7420](#) - Delegate permission checks to FSDirectory
- [HDFS-7463](#) - Simplify FSNamesystem#getBlockLocationsUpdateTimes
- [HDFS-7478](#) - Move org.apache.hadoop.hdfs.server.namenode.NNConf to FSNamesystem
- [HDFS-7517](#) - Remove redundant non-null checks in FSNamesystem#getBlockLocations
- [HDFS-7964](#) - Add support for async edit logging
- [HDFS-8224](#) - Schedule a block for scanning if its metadata file is corrupt
- [HDFS-8269](#) - getBlockLocations() does not resolve the .reserved path and generates incorrect edit logs when updating the atime
- [HDFS-8709](#) - Clarify automatic sync in FSEditLog#logEdit
- [HDFS-9106](#) - Transfer failure during pipeline recovery causes permanent write failures
- [HDFS-9290](#) - DFSClient#callAppend() is not backward compatible for slightly older NameNodes
- [HDFS-9428](#) - Fix intermittent failure of TestDNFencing.testQueueingWithAppend
- [HDFS-9549](#) - TestCacheDirectives#testExceedsCapacity is unreliable
- [HDFS-9630](#) - DistCp minor refactoring and clean up
- [HDFS-9638](#) - Improve DistCp Help and documentation
- [HDFS-9764](#) - DistCp does not print value for several arguments including -numListstatusThreads.
- [HDFS-9781](#) - FsDatasetImpl#getBlockReports can occasionally throw NullPointerException
- [HDFS-9820](#) - Improve distcp to support efficient restore to an earlier snapshot
- [HDFS-9906](#) - Remove excessive unnecessary log output when a DataNode is restarted
- [HDFS-9958](#) - BlockManager#createLocatedBlocks can throw NPE for corruptBlocks on failed storages
- [HDFS-10178](#) - Permanent write failures can happen if pipeline recoveries occur for the first packet
- [HDFS-10216](#) - Distcp -diff throws exception when handling relative path
- [HDFS-10270](#) - TestJMXGet:testNameNode() fails
- [HDFS-10271](#) - Extra bytes are getting released from reservedSpace for append
- [HDFS-10298](#) - Document the usage of distcp -diff option
- [HDFS-10313](#) - Distcp need to enforce the order of snapshot names passed to -diff
- [HDFS-10397](#) - Distcp should ignore -delete option if -diff option is provided instead of exiting
- [HDFS-10457](#) - DataNode should not auto-format block pool directory if VERSION is missing.
- [HDFS-10525](#) - Fix NPE in CacheReplicationMonitor#rescanCachedBlockMap
- [HDFS-10556](#) - DistCpOptions should be validated automatically
- [HDFS-10609](#) - Uncaught InvalidEncryptionKeyException during pipeline recovery can abort downstream applications
- [HDFS-10722](#) - Fix race condition in TestEditLog#testBatchedSyncWithClosedLogs
- [HDFS-10760](#) - DataXceiver#run() should not log InvalidToken exception as an error
- [HDFS-10822](#) - Log DataNodes in the write pipeline. John Zhuge via Lei Xu
- [HDFS-10879](#) - TestEncryptionZonesWithKMS#testReadWrite fails intermittently
- [HDFS-10962](#) - TestRequestHedgingProxyProvider is unreliable
- [HDFS-11012](#) - Unnecessary INFO logging on DFSClients for InvalidToken
- [HDFS-11040](#) - Add documentation for HDFS-9820 distcp improvement
- [HDFS-11056](#) - Concurrent append and read operations lead to checksum error
- [MAPREDUCE-6359](#) - In RM HA setup, Cluster tab links populated with AM hostname instead of RM
- [MAPREDUCE-6473](#) - Job submission can take a long time during Cluster initialization
- [MAPREDUCE-6635](#) - Unsafe long to int conversion in UncompressedSplitLineReader and IndexOutOfBoundsExcpetion

- [MAPREDUCE-6680](#) - JHS UserLogDir scan algorithm sometime could skip directory with update in CloudFS (Azure FileSystem, S3, etc)
- [MAPREDUCE-6684](#) - High contention on scanning of user directory under immediate_done in Job History Server
- [MAPREDUCE-6761](#) - Regression when handling providers - invalid configuration ServiceConfiguration causes Cluster initialization failure
- [MAPREDUCE-6771](#) - RMContainerAllocator sends container diagnostics event after corresponding completion event
- [YARN-3495](#) - Confusing log generated by FairScheduler
- [YARN-4004](#) - container-executor should print output of docker logs if the docker container exits with non-0 exit status
- [YARN-4017](#) - container-executor overuses PATH_MAX
- [YARN-4245](#) - Generalize config file handling in container-executor
- [YARN-4255](#) - container-executor does not clean up Docker operation command files
- [YARN-5608](#) - TestAMRMClient.setup() fails with ArrayOutOfBoundsException
- [YARN-5704](#) - Provide config knobs to control enabling/disabling new/work in progress features in container-executor
- [HBASE-13330](#) - Region left unassigned due to AM & SSH each thinking the assignment would be done by the other
- [HBASE-14241](#) - Fix deadlock during cluster shutdown due to concurrent connection close
- [HBASE-14313](#) - After a Connection sees ConnectionClosingException on a connection it never recovers
- [HBASE-14407](#) - NotServingRegion: hbase region closed forever
- [HBASE-14449](#) - Rewrite deadlock prevention for concurrent connection close
- [HBASE-14474](#) - Addendum closes connection in writeRequest() outside synchronized block
- [HBASE-14474](#) - DeadLock in RpcClientImpl.Connection.close()
- [HBASE-14578](#) - URISyntaxException during snapshot restore for table with user defined namespace
- [HBASE-14968](#) - ConcurrentModificationException in region close resulting in the region staying in closing state
- [HBASE-15430](#) - Failed taking snapshot - Manifest proto-message too large
- [HBASE-15856](#) - Addendum Fix UnknownHostException import in MetaTableLocator
- [HBASE-15856](#) - Do not cache unresolved addresses for connections
- [HBASE-16350](#) - Undo server abort from HBASE-14968
- [HBASE-16360](#) - TableMapReduceUtil addHBaseDependencyJars has the wrong class name for PrefixTreeCodec
- [HBASE-16767](#) - Mob compaction needs to clean up files in /hbase/mobdir/.tmp and /hbase/mobdir/.tmp/.bulkload when running into IO exceptions
- [HIVE-10384](#) - BackportRetryingMetaStoreClient does not retry wrapped TTransportExceptions
- [HIVE-10728](#) - Deprecate unix_timestamp(void) and make it deterministic
- [HIVE-11579](#) - Invoke the set command will close standard error output[beeline-cli]
- [HIVE-11768](#) - java.io.DeleteOnExitHook leaks memory on long running Hive Server2 Instances
- [HIVE-11901](#) - StorageBasedAuthorizationProvider requires write permission on table for SELECT statements
- [HIVE-12475](#) - Parquet schema evolution within array<struct<>> does not work
- [HIVE-12891](#) - Hive fails when java.io.tmpdir is set to a relative location
- [HIVE-13058](#) - Add session and operation_log directory deletion messages
- [HIVE-13090](#) - Hive metastore crashes on NPE with ZooKeeperTokenStore
- [HIVE-13129](#) - CliService leaks HMS connection
- [HIVE-13198](#) - Authorization issues with cascading views
- [HIVE-13237](#) - Select parquet struct field with upper case throws NPE
- [HIVE-13429](#) - Tool to remove dangling scratch directory
- [HIVE-13997](#) - Insert overwrite directory does not overwrite existing files
- [HIVE-14296](#) - Session count is not decremented when HS2 clients do not shutdown cleanly
- [HIVE-14421](#) - FS.deleteOnExit holds references to _tmp_space.db files
- [HIVE-14436](#) - Hive 1.2.1/Hitting "ql.Driver: FAILED: IllegalArgumentException Error: , expected at the end of 'decimal(9'" after enabling hive.optimize.skewjoin and with MR engine
- [HIVE-14457](#) - Partitions in encryption zone are still trashed though an exception is returned
- [HIVE-14743](#) - ArrayIndexOutOfBoundsException - HBASE-backed views' query with JOINS

- [HIVE-14762](#) - Add logging while removing scratch space
- [HIVE-14805](#) - Subquery inside a view will have the object in the subquery as the direct input
- [HIVE-14817](#) - Shut down the SessionManager timeoutChecker thread properly upon shutdown
- [HIVE-15090](#) - Temporary DB failure can stop ExpiredTokenRemover thread
- [HUE-4804](#) - [search] Download function of HTML widget breaks the display
- [HUE-4968](#) - [oozie] Remove access to /oozie/import_workflow when v2 is enabled
- [IMPALA-1928](#) - Fix Thrift client transport wrapping order
- [IMPALA-3369](#) - Add ALTER TABLE SET COLUMN STATS statement
- [IMPALA-3378](#) - HiveUdfCall::Open() produces unsynchronized access to JNIUtil::global_refs_vector
- [IMPALA-3379](#) - HBaseTableWriter::CreatePutList() produces unsynchronized access to JNIUtil::global_refs_vector
- [IMPALA-3441](#) - Check for malformed Avro data
- [IMPALA-3499](#) - Split catalog update
- [IMPALA-3575](#) - Add retry to backend connection request and rpc timeout
- [IMPALA-3633](#) - Cancel fragment if coordinator is gone
- [IMPALA-3682](#) - Do not retry unrecoverable socket creation errors
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3698](#) - Fix Isilon permissions test
- [IMPALA-3711](#) - Remove unnecessary privilege checks in getDbsMetadata()
- [IMPALA-3732](#) - Handle string length overflow in Avro files
- [IMPALA-3751](#) - Fix clang build errors and warnings
- [IMPALA-3915](#) - Register privilege and audit requests when analyzing resolved table refs.
- [IMPALA-4135](#) - Thrift threaded server times-out connections during high load
- [IMPALA-4153](#) - Return valid non-NULL pointer for 0-byte allocations
- [OOZIE-1814](#) - Oozie should mask any passwords in logs and REST interfaces
- [OOZIE-2068](#) - Configuration as part of sharelib
- [OOZIE-2347](#) - Remove unnecessary new Configuration()/new jobConf() calls from Oozie
- [OOZIE-2555](#) - Oozie SSL enable setup does not return port for admin -servers
- [OOZIE-2567](#) - HCat connection is not closed while getting hcat cred
- [OOZIE-2589](#) - CompletedActionXCommand is hardcoded to wrong priority
- [OOZIE-2649](#) - Cannot override sub-workflow configuration property if defined in parent workflow XML
- [PIG-3569](#) - SUM function for BigDecimal and BigInteger
- [PIG-3818](#) - PIG-2499 is accidentally reverted
- [SENTRY-1095](#) - Insert into requires URI privilege on partition location under table.
- [SOLR-7280](#) - BackportLoad cores in sorted order and tweak coreLoadThread counts to improve cluster stability on restarts
- [SOLR-8407](#) - Backport
- [SOLR-8586](#) - Add index fingerprinting and use it in peersync
- [SOLR-8690](#) - Add solr.disableFingerprint system property
- [SOLR-8691](#) - Cache index fingerprints per searcher
- [SOLR-9310](#) - SOLR-9524
- [SPARK-17644](#) - [CORE] Do not add failedStages when abortStage for fetch failure
- [SQOOP-2387](#) - Sqoop should support importing from table with column names containing some special character
- [SQOOP-2864](#) - ClassWriter chokes on column names containing double quotes
- [SQOOP-2880](#) - Provide argument for overriding temporary directory
- [SQOOP-2884](#) - Document --temporary-rootdir
- [SQOOP-2906](#) - Optimization of AvroUtil.toAvroIdentifier
- [SQOOP-2915](#) - Fixing Oracle related unit tests
- [SQOOP-2920](#) - Sqoop performance deteriorates significantly on wide datasets; Sqoop 100% on CPU
- [SQOOP-2952](#) - Fixing bug row key not added into column family using --hbase-bulkload
- [SQOOP-2971](#) - OraOop does not close connections properly

- [SQOOP-2983](#) - OraOop export has degraded performance with wide tables
- [SQOOP-2986](#) - Add validation check for --hive-import and --incremental lastmodified
- [SQOOP-3021](#) - ClassWriter fails if a column name contains a backslash character
- [SQOOP-3034](#) - HBase import should fail fast if using anything other than as-textfile

Issues Fixed in CDH 5.5.5

CDH 5.5.5 fixes the following issues.

Apache Oozie

Oozie Web Console returns 500 error when Oozie server runs on JDK 8u75 or higher

Bug: [OOZIE-2533](#)

Cloudera Bug: CDH-40362

The Oozie Web Console returns a 500 error when the Oozie server is running on JDK 8u75 and higher. The Oozie server still functions, and you can use the Oozie command line, REST API, Java API, or the Hue Oozie Dashboard to review status of those jobs.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.5.5:

- [FLUME-2821](#) - KafkaSourceUtil Can Log Passwords at Info remove logging of security related data in older releases
- [FLUME-2913](#) - Don't strip SLF4J from imported classpaths
- [FLUME-2918](#) - Speed up TaildirSource on directories with many files
- [HADOOP-8436](#) - NPE In getLocalPathForWrite (path, conf) when the required context item is not configured
- [HADOOP-8437](#) - getLocalPathForWrite should throw IOException for invalid paths
- [HADOOP-8751](#) - NPE in Token.toString() when Token is constructed using null identifier
- [HADOOP-8934](#) - Shell command ls should include sort options
- [HADOOP-10048](#) - LocalDirAllocator should avoid holding locks while accessing the filesystem
- [HADOOP-10971](#) - Add -C flag to make `hadoop fs -ls` print filenames only
- [HADOOP-11901](#) - BytesWritable fails to support 2G chunks due to integer overflow
- [HADOOP-11984](#) - Enable parallel JUnit tests in pre-commit
- [HADOOP-12252](#) - LocalDirAllocator should not throw NPE with empty string configuration
- [HADOOP-12259](#) - Utility to Dynamic port allocation
- [HADOOP-12659](#) - Incorrect usage of config parameters in token manager of KMS
- [HADOOP-12787](#) - KMS SPNEGO sequence does not work with WEBHDFS
- [HADOOP-12841](#) - Update s3-related properties in core-default.xml.
- [HADOOP-12901](#) - Add warning log when KMSClientProvider cannot create a connection to the KMS server.
- [HADOOP-12963](#) - Allow using path style addressing for accessing the s3 endpoint.
- [HADOOP-13079](#) - Add -q option to Ls to print ? instead of non-printable characters
- [HADOOP-13132](#) - Handle ClassCastException on AuthenticationException in LoadBalancingKMSClientProvider
- [HADOOP-13155](#) - Implement TokenRenewer to renew and cancel delegation tokens in KMS
- [HADOOP-13251](#) - Authenticate with Kerberos credentials when renewing KMS delegation token
- [HADOOP-13255](#) - KMSClientProvider should check and renew tgt when doing delegation token operations
- [HADOOP-13263](#) - Reload cached groups in background after expiry.
- [HADOOP-13457](#) - Remove hardcoded absolute path for shell executable.
- [HDFS-6434](#) - Default permission for creating file should be 644 for WebHdfs/HttpFS
- [HDFS-7597](#) - DelegationTokenIdentifier should cache the TokenIdentifier to UGI mapping
- [HDFS-8008](#) - Support client-side back off when the datanodes are congested
- [HDFS-8581](#) - ContentSummary on / skips further counts on yielding lock
- [HDFS-8829](#) - Make SO_RCVBUF and SO_SNDBUF size configurable for DataTransferProtocol sockets and allow configuring auto-tuning
- [HDFS-8897](#) - Balancer should handle fs.defaultFS trailing slash in HA

- [HDFS-9085](#) - Show renewer information in DelegationTokenIdentifier#toString
- [HDFS-9259](#) - Make SO_SNDBUF size configurable at DFSCliet side for hdfs write scenario.
- [HDFS-9276](#) - Failed to Update HDFS Delegation Token for long running application in HA mode
- [HDFS-9365](#) - Balancer does not work with the HDFS-6376 HA setup.
- [HDFS-9405](#) - Warmup NameNode EDEK caches in background thread
- [HDFS-9466](#) - TestShortCircuitCache#testDataXceiverCleansUpSlotsOnFailure is flaky
- [HDFS-9700](#) - DFSCliet and DFSOutputStream should set TCP_NODELAY on sockets for DataTransferProtocol
- [HDFS-9732](#) - , Improve DelegationTokenIdentifier.toString() for better logging
- [HDFS-9805](#) - Add server-side configuration for enabling TCP_NODELAY for DataTransferProtocol and default it to true
- [HDFS-9939](#) - Increase DecompressorStream skip buffer size
- [HDFS-10360](#) - DataNode may format directory and lose blocks if current/VERSION is missing.
- [HDFS-10381](#) - DataStreamer DataNode exclusion log message should be warning.
- [HDFS-10396](#) - Using -diff option with DistCp may get "Comparison method violates its general contract" exception
- [HDFS-10481](#) - HTTPFS server should correctly impersonate as end user to open file
- [HDFS-10512](#) - VolumeScanner may terminate due to NPE in DataNode.reportBadBlocks
- [HDFS-10516](#) - Fix bug when warming up EDEK cache of more than one encryption zone
- [HDFS-10544](#) - Balancer doesn't work with IPFailoverProxyProvider.
- [HDFS-10643](#) - Namenode should use loginUser(hdfs) to generateEncryptedKey
- [MAPREDUCE-6442](#) - Stack trace is missing when error occurs in client protocol provider's constructor
- [MAPREDUCE-6473](#) - Job submission can take a long time during Cluster initialization
- [MAPREDUCE-6577](#) - MR AM unable to load native library without MR_AM_ADMIN_USER_ENV set
- [YARN-2605](#) - [RM HA] Rest api endpoints doing redirect incorrectly.
- [YARN-3055](#) - Fixed ResourceManager's DelegationTokenRenewer to not stop token renewal of applications part of a bigger workflow
- [YARN-3104](#) - Fixed RM to not generate new AMRM tokens on every heartbeat between rolling and activation
- [YARN-3832](#) - Resource Localization fails on a cluster due to existing cache directories
- [YARN-4459](#) - container-executor should only kill process groups
- [YARN-4784](#) - Fairscheduler: defaultQueueSchedulingPolicy should not accept FIFO.
- [YARN-5048](#) - DelegationTokenRenewer#skipTokenRenewal may throw NPE
- [YARN-5272](#) - Handle queue names consistently in FairScheduler.
- [HBASE-11625](#) - Verifies data before building HFileBlock.
- [HBASE-14155](#) - StackOverflowError in reverse scan
- [HBASE-14644](#) - Region in transition metric is broken
- [HBASE-14730](#) - Region server needs to log warnings when there are attributes configured for cells with hfile v2
- [HBASE-15439](#) - getMaximumAllowedTimeBetweenRuns in ScheduledChore ignores the TimeUnit
- [HBASE-15496](#) - Throw RowTooBigException only for user scan/get
- [HBASE-15707](#) - ImportTSV bulk output does not support tags with hfile.format.version=3
- [HBASE-15746](#) - Remove extra RegionCoprocesor preClose() in RSRpcServices#closeRegion
- [HBASE-15791](#) - Improve javadoc around ScheduledChore
- [HBASE-15811](#) - Batch Get after batch Put does not fetch all Cells
- [HBASE-15925](#) - Provide default values for hadoop compat module related properties that match default hadoop profile.
- [HBASE-16207](#) - Can't restore snapshot without "Admin" permission
- [HBASE-16288](#) - Revert "HFile intermediate block level indexes might recurse forever creating multi TB files"
- [HBASE-16288](#) - HFile intermediate block level indexes might recurse forever creating multi TB files
- [HIVE-7443](#) - Fix HiveConnection to communicate with Kerberized Hive JDBC server and alternative JDKs
- [HIVE-9499](#) - hive.limit.query.max.table.partition makes queries fail on non-partitioned tables
- [HIVE-10685](#) - Alter table concatenate oparetor will cause duplicate data
- [HIVE-10925](#) - Non-static threadlocals in metastore code can potentially cause memory leak

- [HIVE-11031](#) - ORC concatenation of old files can fail while merging column statistics
- [HIVE-11243](#) - Changing log level in Utilities.getBaseWork
- [HIVE-11369](#) - Mapjoins in HiveServer2 fail when jmxremote is used
- [HIVE-11408](#) - HiveServer2 is leaking ClassLoaders when add jar / temporary functions are used due to constructor caching in Hadoop ReflectionUtils
- [HIVE-11427](#) - Location of temporary table for CREATE TABLE SELECT broken by HIVE-7079.
- [HIVE-11747](#) - Unnecessary error log is shown when executing a "INSERT OVERWRITE LOCAL DIRECTORY" cmd in the embedded mode
- [HIVE-11827](#) - STORED AS AVRO fails SELECT COUNT(*) when empty
- [HIVE-12481](#) - Occasionally "Request is a replay" will be thrown from HS2
- [HIVE-12635](#) - Hive should return the latest hbase cell timestamp as the row timestamp value
- [HIVE-12958](#) - Make embedded Jetty server more configurable
- [HIVE-13285](#) - Orc concatenation may drop old files from moving to final path
- [HIVE-13462](#) - HiveResultSetMetaData.getPrecision() fails for NULL columns
- [HIVE-13527](#) - Using deprecated APIs in HBase client causes zookeeper connection leaks
- [HIVE-13570](#) - Some queries with Union all fail when CBO is off
- [HIVE-13590](#) - Kerberized HS2 with LDAP auth enabled fails in multi-domain LDAP case
- [HIVE-13704](#) - Don't call DistCp.execute() instead of DistCp.run()
- [HIVE-13736](#) - View's input/output formats are TEXT by default.
- [HIVE-13932](#) - Hive SMB Map Join with small set of LIMIT failed with NPE
- [HIVE-13953](#) - Issues in HiveLockObject equals method
- [HIVE-13991](#) - Union All on view fail with no valid permission on underneath table
- [HIVE-14006](#) - Hive query with UNION ALL fails with ArrayIndexOutOfBoundsException.
- [HIVE-14118](#) - Make the alter partition exception more meaningful
- [HUE-3520](#) - [jib] Fix backport error
- [HUE-3520](#) - [jib] Use impersonation to access JHS if security is enabled
- [HUE-3637](#) - [sqoop] Avoid decode errors on attribute values
- [HUE-3650](#) - [beeswax] Notify of caught errors in the watch logs process
- [HUE-3651](#) - [core] Upgrade Moment.js
- [HUE-3716](#) - [core] Add gen-py paths to hue.pth
- [HUE-3861](#) - [core] Upgrade Django Axes to 1.5
- [HUE-3866](#) - [core] Hue CPU reaches ~100% usage while uploading files with SSL to HTTPFS/WebHDFS
- [HUE-3880](#) - [core] Add importlib directly for Python 2.6
- [HUE-4005](#) - [oozie] Remove oozie.coord.application.path from properties when rerunning workflow
- [HUE-4006](#) - [oozie] Create new deployment directory when coordinator or bundle is copied
- [HUE-4007](#) - [oozie] Fix deployment_dir for the bundle in oozie example fixtures
- [HUE-4023](#) - [useradmin] update AuthenticationForm to allow activated users to login
- [HUE-4087](#) - [jobbrowser] Unable to kill jobs with Resource Manager HA enabled
- [HUE-4202](#) - [jib] Enable offset param for fetching jobbrowser logs
- [HUE-4215](#) - [yarn] Reset API_CACHE on logout
- [HUE-4227](#) - [yarn] Fix unittest for MR API Cache
- [HUE-4238](#) - [doc2] Ignore history docs in find_jobs_with_no_doc during sync documents
- [HUE-4252](#) - [core] Handle 307 redirect from YARN upon standby failover
- [HUE-4258](#) - [jib] Close and pool Spark History Server connections
- [HUE-4333](#) - [core] Properly reset API_CACHE on failover
- [HUE-4493](#) - [oozie] Fix sync-workflow action when Workflow includes sub-workflow
- [HUE-4515](#) - [oozie] Remove oozie.bundle.application.path from properties when rerunning workflow
- [OOZIE-2314](#) - Unable to kill old instance child job by workflow or coord rerun by Launcher
- [OOZIE-2329](#) - Make handling yarn restarts configurable

- [OOZIE-2330](#) - Spark action should take the global jobTracker and nameNode configs by default and allow file and archive elements
- [OOZIE-2345](#) - Parallel job submission for forked actions
- [OOZIE-2391](#) - spark-opts value in workflow.xml is not parsed properly
- [OOZIE-2436](#) - Fork/join workflow fails with oozie.action.yarn.tag must not be null
- [OOZIE-2481](#) - Add YARN_CONF_DIR in the Shell action
- [OOZIE-2504](#) - Create a log4j.properties under HADOOP_CONF_DIR in Shell Action
- [OOZIE-2511](#) - SubWorkflow missing variable set from option if config-default is present in parent workflow
- [OOZIE-2533](#) - Patch-1550 - workaround
- [OOZIE-2537](#) - SqoopMain does not set up log4j properly
- [SENTRY-1190](#) - IMPORT TABLE silently fails if Sentry is enabled
- [SENTRY-1201](#) - Sentry ignores database prefix for MSCK statement
- [SENTRY-1252](#) - grantServerPrivilege and revokeServerPrivilege should treat "*" and "ALL" as synonyms when action is not explicitly specified
- [SENTRY-1265](#) - Sentry service should not require a TGT as it is not talking to other kerberos services as a client
- [SENTRY-1292](#) - Reorder DBModelAction EnumSet
- [SENTRY-1293](#) - Avoid converting string permission to Privilege object
- [SOLR-6631](#) - DistributedQueue spinning on calling zookeeper getChildren()
- [SOLR-6879](#) - Have an option to disable autoAddReplicas temporarily for all collections.
- [SOLR-7178](#) - OverseerAutoReplicaFailoverThread compares Integer objects using ==
- [SOLR-8451](#) - Fix backport
- [SOLR-8497](#) - Merge indexes should mark its directories as done rather than keep them around in the directory cache.
- [SOLR-8551](#) - Make collection deletion more robust.
- [SOLR-8683](#) - Tune down stream closed logging
- [SOLR-9236](#) - AutoAddReplicas will append an extra /tlog to the update log location on replica failover.
- [SPARK-10577](#) - [PYSARK] DataFrame hint for broadcast join
- [SPARK-11442](#) - Reduce numSlices for local metrics test of SparkListenerSuite
- [SPARK-12087](#) - [STREAMING] Create new JobConf for every batch in saveAsHadoopFiles
- [SQOOP-2846](#) - Sqoop Export with update-key failing for avro data file

Issues Fixed in CDH 5.5.4

CDH 5.5.4 fixes the following issues.

Apache Hadoop

FSImage may get corrupted after deleting snapshot

Bug: [HDFS-9406](#)

Cloudera Bug: CDH-33224

When deleting a snapshot that contains the last record of a given INode, the fsimage may become corrupt because the create list of the snapshot diff in the previous snapshot and the child list of the parent INodeDirectory are not cleaned.

Apache HBase

The ReplicationCleaner process can abort if its connection to ZooKeeper is inconsistent

Bug: [HBASE-15234](#)

If the connection with ZooKeeper is inconsistent, the ReplicationCleaner may abort, and the following event is logged by the HMaster:

```
WARN org.apache.hadoop.hbase.replication.master.ReplicationLogCleaner: Aborting
ReplicationLogCleaner
because Failed to get list of replicators
```

Unprocessed WALs accumulate.

The `seekBefore()` method calculates the size of the previous data block by assuming that data blocks are contiguous, and HFile v2 and higher store Bloom blocks and leaf-level INode blocks with the data. As a result, reverse scans do not work when Bloom blocks or leaf-level INode blocks are present when HFile v2 or higher is used.

Workaround: Restart the HMaster occasionally. The `ReplicationCleaner` restarts if necessary and process the unprocessed WALs.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.5.4:

- [FLUME-2632](#) - High CPU on KafkaSink
- [FLUME-2712](#) - Optional channel errors slows down the Source to Main channel event rate
- [FLUME-2781](#) - Kafka Channel with `parseAsFlumeEvent=true` should write data as is, not as flume events
- [FLUME-2886](#) - Optional Channels can cause OOMs
- [FLUME-2891](#) - Revert FLUME-2712 and FLUME-2886
- [FLUME-2897](#) - AsyncHBase sink NPE when `Channel.getTransaction()` fails
- [HADOOP-7139](#) - Allow appending to existing SequenceFiles
- [HADOOP-7817](#) - `RawLocalFileSystem.append()` should give `FSDDataOutputStream` with accurate `.getPos()`
- [HADOOP-11321](#) - `copyToLocal` cannot save a file to an SMB share unless the user has Full Control permissions
- [HADOOP-11687](#) - Ignore `x-*` and response headers when copying an Amazon S3 object
- [HADOOP-11722](#) - Some Instances of Services using `ZKDelegationTokenSecretManager` go down when old token cannot be deleted
- [HADOOP-12240](#) - Fix tests requiring native library to be skipped in non-native profile
- [HADOOP-12280](#) - Skip unit tests based on maven profile rather than `NativeCodeLoader.isNativeCodeLoaded`
- [HADOOP-12559](#) - KMS connection failures should trigger TGT renewal
- [HADOOP-12605](#) - Fix intermittent failure of `TestIPC.testIpcWithReaderQueuing`
- [HADOOP-12668](#) - Support excluding weak Ciphers in `HttpServer2` through `ssl-server.conf`
- [HADOOP-12682](#) - Fix `TestKMS#testKMSRestart*` failure
- [HADOOP-12699](#) - `TestKMS#testKMSProvider` intermittently fails during 'test rollover draining'
- [HADOOP-12715](#) - `TestValueQueue#testgetAtMostPolicyALL` fails intermittently
- [HADOOP-12718](#) - Incorrect error message by `fs -put local dir without permission`
- [HADOOP-12736](#) - `TestTimedOutTestsListener#testThreadDumpAndDeadlocks` sometimes times out
- [HADOOP-12788](#) - `OpensslAesCtrCryptoCodec` should log which random number generator is used
- [HADOOP-12825](#) - Log slow name resolutions
- [HADOOP-12954](#) - Add a way to change `hadoop.security.token.service.use_ip`
- [HADOOP-12972](#) - `Lz4Compressor#getLibraryName` returns the wrong version number
- [HDFS-6520](#) - `hdfs fsck` passes invalid length value when creating `BlockReader`
- [HDFS-7373](#) - Clean up temporary files after `fsimage` transfer failures
- [HDFS-7758](#) - Retire `FsDatasetSpi#getVolumes()` and use `FsDatasetSpi#getVolumeRefs()` instead
- [HDFS-8211](#) - `DataNode` UUID is always null in the JMX counter
- [HDFS-8496](#) - Calling `stopWriter()` with `FSDatasetImpl` lock held may block other threads
- [HDFS-8576](#) - Lease recovery should return true if the lease can be released and the file can be closed
- [HDFS-8785](#) - `TestDistributedFileSystem` is failing in trunk
- [HDFS-8855](#) - `Webhdfs` client leaks active `NameNode` connections
- [HDFS-9264](#) - Minor cleanup of operations on `FsVolumeList#volumes`
- [HDFS-9289](#) - Make `DataStreamer#block` thread safe and verify `genStamp` in `commitBlock`
- [HDFS-9347](#) - Invariant assumption in `TestQuorumJournalManager.shutdown()` is wrong
- [HDFS-9350](#) - Avoid creating temporary strings in `Block.toString()` and `getBlockName()`
- [HDFS-9358](#) - `TestNodeCount#testNodeCount` timed out
- [HDFS-9406](#) - `FSImage` may get corrupted after deleting snapshot
- [HDFS-9514](#) - `TestDistributedFileSystem.testDFSClientPeerWriteTimeout` failing; exception being swallowed

- [HDFS-9576](#) - HTrace: collect position/length information on read operations
- [HDFS-9589](#) - Block files which have been hardlinked should be duplicated before the DataNode appends to the them
- [HDFS-9612](#) - DistCp worker threads are not terminated after jobs are done
- [HDFS-9655](#) - NN should start JVM pause monitor before loading fsimage.
- [HDFS-9688](#) - Test the effect of nested encryption zones in HDFS downgrade
- [HDFS-9701](#) - DN may deadlock when hot-swapping under load
- [HDFS-9721](#) - Allow Delimited PB OIV tool to run upon fsimage that contains INodeReference
- [HDFS-9949](#) - Add a test case to ensure that the DataNode does not regenerate its UUID when a storage directory is cleared
- [HDFS-10223](#) - peerFromSocketAndKey performs SASL exchange before setting connection timeouts
- [HDFS-10267](#) - Extra "synchronized" on FsDatasetImpl#recoverAppend and FsDatasetImpl#recoverClose
- [MAPREDUCE-4785](#) - TestMRApp occasionally fails
- [MAPREDUCE-6460](#) - TestRMContainerAllocator.testAttemptNotFoundCausesRMCommunicatorException fails
- [MAPREDUCE-6528](#) - Memory leak for HistoryFileManager.getJobSummary()
- [MAPREDUCE-6580](#) - Test failure: TestMRJobsWithProfiler
- [MAPREDUCE-6620](#) - Jobs that did not start are shown as starting in 1969 in the JHS web UI
- [YARN-2749](#) - Fix some testcases from TestLogAggregationService fails in trunk
- [YARN-2871](#) - TestRMRestart#testRMRestartGetApplicationList sometime fails in trunk
- [YARN-2902](#) - Killing a container that is localizing can orphan resources in the DOWNLOADING state
- [YARN-3446](#) - FairScheduler headroom calculation should exclude nodes in the blacklist
- [YARN-3727](#) - For better error recovery, check if the directory exists before using it for localization
- [YARN-4155](#) - TestLogAggregationService.testLogAggregationServiceWithInterval failing
- [YARN-4168](#) - Fixed a failing test TestLogAggregationService.testLocalFileDeletionOnDiskFull
- [YARN-4354](#) - Public resource localization fails with NPE
- [YARN-4380](#) - TestResourceLocalizationService.testDownloadingResourcesOnContainerKill fails intermittently
- [YARN-4393](#) - Fix intermittent test failure for TestResourceLocalizationService#testFailedDirsResourceRelease
- [YARN-4546](#) - ResourceManager crash due to scheduling opportunity overflow
- [YARN-4573](#) - Fix test failure in TestRMAppTransitions#testAppRunningKill and testAppKilledKilled
- [YARN-4613](#) - Fix test failure in TestClientRMService#testGetClusterNodes
- [YARN-4704](#) - TestResourceManager#testResourceAllocation() fails when using FairScheduler
- [YARN-4717](#) - TestResourceLocalizationService.testPublicResourceInitializesLocalDir fails Intermittently due to IllegalArgumentException from cleanup
- [HBASE-6617](#) - ReplicationSourceManager should be able to track multiple WAL paths (ADDENDUM)
- [HBASE-14586](#) - Use a maven profile to run Jacoco analysis
- [HBASE-14587](#) - Attach a test-sources.jar for hbase-server
- [HBASE-14588](#) - Stop accessing test resources from within src folder
- [HBASE-14759](#) - Avoid using Math.abs when selecting SyncRunner in FSHLog
- [HBASE-15019](#) - Replication stuck when HDFS is restarted
- [HBASE-15052](#) - Use EnvironmentEdgeManager in ReplicationSource
- [HBASE-15152](#) - Automatically include prefix-tree module in MR jobs if present
- [HBASE-15157](#) - Add *PerformanceTest for Append, CheckAnd*
- [HBASE-15206](#) - Fix flaky testSplitDaughtersNotInMeta
- [HBASE-15213](#) - Fix increment performance regression caused by HBASE-8763 on branch-1.0
- [HBASE-15234](#) - Don't abort ReplicationLogCleaner on ZooKeeper errors
- [HBASE-15456](#) - CreateTableProcedure/ModifyTableProcedure needs to fail when there is no family in table descriptor
- [HBASE-15479](#) - No more garbage or beware of autoboxing
- [HBASE-15582](#) - SnapshotManifestV1 too verbose when there are no regions
- [HIVE-9617](#) - UDF from_utc_timestamp throws NPE if the second argument is null

- [HIVE-9743](#) - Revert "(Tests portion only)Incorrect result set for vectorized left outer join (Matt McCline, reviewed by Vikram Dixit)"
- [HIVE-10115](#) - HS2 running on a Kerberized cluster should offer Kerberos(GSSAPI) and Delegation token(DIGEST) when alternate authentication is enabled
- [HIVE-10213](#) - MapReduce jobs using dynamic-partitioning fail on commit
- [HIVE-10303](#) - HIVE-9471 broke forward compatibility of ORC files
- [HIVE-11054](#) - Handle varchar/char partition columns in vectorization
- [HIVE-11097](#) - HiveInputFormat uses String.startsWith to compare splitPath and PathToAliases
- [HIVE-11135](#) - Fix the Beeline set and save command in order to avoid the NullPointerException
- [HIVE-11285](#) - ObjectInspector for partition columns in FetchOperator in SMBJoin causes exception
- [HIVE-11488](#) - Need to add support for sessionId and queryId logging, QueryId can't be stored in the configuration of the SessionState since multiple queries can run in a single session
- [HIVE-11583](#) - When PTF is used over a large partitions result could be corrupted
- [HIVE-11590](#) - AvroDeserializer is very chatty
- [HIVE-11828](#) - beeline -f fails on scripts with tabs between column type and comment
- [HIVE-11866](#) - Add framework to enable testing using LDAPServer using LDAP protocol
- [HIVE-11919](#) - Hive Union Type Mismatch
- [HIVE-12315](#) - Fix Vectorized double divide by zero
- [HIVE-12354](#) - MapJoin with double keys is slow on MR
- [HIVE-12431](#) - Support timeout for compile lock
- [HIVE-12506](#) - SHOW CREATE TABLE command creates a table that does not work for RCFile format
- [HIVE-12706](#) - Incorrect output from from_utc_timestamp()/to_utc_timestamp when local timezone has DST
- [HIVE-12782](#) - Update the golden files for some tests that fail
- [HIVE-12790](#) - Metastore connection leaks in HiveServer2
- [HIVE-12885](#) - LDAP Authenticator improvements
- [HIVE-12909](#) - Some encryption q-tests fail because trash is disabled in encryption_with_trash.q
- [HIVE-12941](#) - Unexpected result when using MIN() on struct with NULL in first field
- [HIVE-12946](#) - Alter table should also add default scheme and authority for the location similar to create table
- [HIVE-13039](#) - BETWEEN predicate is not functioning correctly with predicate pushdown on Parquet table
- [HIVE-13055](#) - Add unit tests for HIVE-11512
- [HIVE-13065](#) - Hive throws NPE when writing map type data to a HBase backed table
- [HIVE-13082](#) - Enable constant propagation optimization in query with left semi join
- [HIVE-13200](#) - Aggregation functions returning empty rows on partitioned columns
- [HIVE-13243](#) - Hive drop table on encryption zone fails for external tables
- [HIVE-13251](#) - Hive can't read the decimal in AVRO file generated from previous version
- [HIVE-13286](#) - Query ID is being reused across queries
- [HIVE-13295](#) - Improvement to LDAP search queries in HS2 LDAP Authenticator
- [HIVE-13401](#) - Kerberized HS2 with LDAP auth enabled fails kerberos/delegation token authentication
- [HUE-3106](#) - [filebrowser] Add support for full paths in zip file uploads
- [HUE-3110](#) - [oozie] Fix bundle submission when coordinator points to multiple bundles
- [HUE-3132](#) - [core] Fix Sync Ldap users and groups for anonymous binds
- [HUE-3180](#) - [useradmin] Override duplicate username validation message
- [HUE-3185](#) - [oozie] Avoid extra API calls for parent information in workflow dashboard
- [HUE-3303](#) - [core] PostgreSQL requires data update and alter table operations in separate transactions
- [HUE-3310](#) - [jobsub] Prevent browsing job designs by API
- [HUE-3334](#) - [editor] Skip checking for multi queries if there is no semi colon, send empty query instead of error
- [HUE-3398](#) - [beeswax] Filter out sessions with empty guid or secret key
- [HUE-3436](#) - [oozie] Retain old dependencies when saving a workflow
- [HUE-3437](#) - [core] PamBackend does not honor ignore_username_case
- [HUE-3523](#) - [oozie] Modify find_jobs_with_no_doc method to exclude jobs with no name

- [HUE-3528](#) - [oozie] Call correct metrics api to avoid 500 error
- [HUE-3594](#) - [fb] Smarter DOM based XSS filter on hashes
- [IMPALA-852](#) - ,IMPALA-2215: Analyze HAVING clause before aggregation
- [IMPALA-1092](#) - Fix estimates for trivial coord-only queries
- [IMPALA-1170](#) - Fix URL parsing when path contains '@'
- [IMPALA-1934](#) - Allow shell to retrieve LDAP password from shell cmd
- [IMPALA-2093](#) - Disallow NOT IN aggregate subqueries with a constant lhs expr
- [IMPALA-2184](#) - Don't inline timestamp methods with try/catch blocks in IR
- [IMPALA-2425](#) - Broadcast join hint not enforced when low memory limit is set
- [IMPALA-2503](#) - Add missing String.format() arg in error message
- [IMPALA-2539](#) - Unmark collections slots of empty union operands
- [IMPALA-2554](#) - Change default buffer size for RPC servers and clients
- [IMPALA-2565](#) - Planner tests are flaky due to file size mismatches
- [IMPALA-2592](#) - DataStreamSender::Channel::CloseInternal() does not close the channel on an error
- [IMPALA-2599](#) - Pseudo-random sleep before acquiring kerberos ticket possibly not really pseudo-random
- [IMPALA-2711](#) - Fix memory leak in Rand()
- [IMPALA-2719](#) - test_parquet_max_page_header fails on Isilon
- [IMPALA-2732](#) - Timestamp formats with non-padded values
- [IMPALA-2734](#) - Correlated EXISTS subqueries with HAVING clause return wrong results
- [IMPALA-2742](#) - Avoid unbounded MemPool growth with AcquireData()
- [IMPALA-2749](#) - Fix decimal multiplication overflow
- [IMPALA-2765](#) - Preserve return type of subexpressions substituted in isTrueWithNullSlots()
- [IMPALA-2788](#) - conv(bigint num, int from_base, int to_base) returns wrong result
- [IMPALA-2798](#) - Bring in AVRO-1617 fix and add test case for it
- [IMPALA-2818](#) - Fix cancellation crashes/hangs due to BlockOnWait() race
- [IMPALA-2820](#) - Support unquoted keywords as struct-field names
- [IMPALA-2832](#) - Fix cloning of FunctionCallExpr
- [IMPALA-2844](#) - Allow count(*) on RC files with complex types
- [IMPALA-2870](#) - Fix failing metadata.test_ddl.TestDdlStatements.test_create_table test
- [IMPALA-2894](#) - Move regression test into a different .test file
- [IMPALA-2906](#) - Fix an edge case with materializing TupleIsNullPredicates in analytic sorts
- [IMPALA-2914](#) - Fix DCHECK Check failed: HasDateOrTime()
- [IMPALA-2926](#) - Fix off-by-one bug in SelectNode::CopyRows()
- [IMPALA-2940](#) - Fix leak of dictionaries in Parquet scanner
- [IMPALA-3000](#) - Fix BitReader::Reset()
- [IMPALA-3034](#) - Verify all consumed memory of a MemTracker is always released at destruction time
- [IMPALA-3047](#) - Separate create table test with nested types
- [IMPALA-3054](#) - Disable probe side filters when spilling
- [IMPALA-3071](#) - Fix assignment of On-clause predicates belonging to an inner join
- [IMPALA-3085](#) - Unregister data sinks' MemTrackers at their Close() functions
- [IMPALA-3093](#) - ReopenClient() could NULL out 'client_key' causing a crash
- [IMPALA-3095](#) - Add configurable whitelist of authorized internal principals
- [IMPALA-3151](#) - Impala crash for avro table when casting to char data type
- [IMPALA-3194](#) - Allow queries materializing scalar type columns in RC/sequence files
- [KITE-1114](#) - Kite CLI json-import HDFS temp file path not multiuser safe, fix missing license header
- [OOZIE-2419](#) - HBase credentials are not correctly proxied
- [OOZIE-2428](#) - TestSLAService, TestSLAEventGeneration flaky tests
- [OOZIE-2429](#) - TestEventGeneration test is flaky
- [OOZIE-2432](#) - TestPurgeXCommand fails
- [OOZIE-2435](#) - TestCoordChangeXCommand is flaky

- [OOZIE-2466](#) - Repeated failure of TestMetricsInstrumentation.testSamplers
- [OOZIE-2486](#) - TestSLAEventsGetForFilterJPAExecutor is flaky
- [OOZIE-2490](#) - Oozie can't set hadoop.security.token.service.use_ip
- [SENTRY-922](#) - BackportINSERT OVERWRITE DIRECTORY permission not working correctly
- [SENTRY-972](#) - backportInclude sentry-tests-hive hadoop test script in maven project
- [SENTRY-991](#) - backportRoles of Sentry Permission needs to be case insensitive
- [SENTRY-1002](#) - PathsUpdate.parsePath(path) will throw an NPE when parsing relative paths
- [SENTRY-1003](#) - Support "reload" by updating the classpath of Sentry function aux jar path during runtime
- [SENTRY-1007](#) - backportSentry column-level performance for wide tables
- [SENTRY-1008](#) - Path should be not be updated if the create/drop table/partition event fails
- [SENTRY-1015](#) - backportImprove Sentry + Hive error message when user has insufficient privileges
- [SENTRY-1044](#) - Tables with non-hdfs locations breaks HMS startup
- [SENTRY-1169](#) - MetastorePlugin#renameAuthzObject log message prints oldpathname as newpathname
- [SENTRY-1184](#) - Clean up HMSPaths.renameAuthzObject
- [SOLR-6820](#) - Make the number of version buckets used by the UpdateLog configurable as increasing beyond the default 256 has been shown to help with high volume indexing performance in SolrCloudIncrease the default number of buckets to 65536 instead of 256, fix numVersionBuckets name attribute in configsets
- [SOLR-7281](#) - Add an overseer action to publish an entire node as 'down'
- [SOLR-7332](#) - Initialize the highest value for all version buckets with the max value from the index or recent updates to avoid unnecessary lookups to the index to check for reordered updates when processing new documents
- [SOLR-7493](#) - Requests aren't distributed evenly if the collection isn't present locally. Merges r1683946 and r1683948 from trunk
- [SOLR-7587](#) - TestSpellCheckResponse stalled and never timed out -- possible VersionBucket bug?
- [SOLR-7625](#) - Version bucket seed not updated after new index is installed on a replica
- [SOLR-8215](#) - Only active replicas should handle incoming requests against a collection
- [SOLR-8371](#) - Try and prevent too many recovery requests from stacking up and clean up some faulty cancel recovery logic
- [SOLR-8451](#) - We should not call method.abort in HttpSolrClient or HttpSolrCall#remoteQuery and HttpSolrCall#remoteQuery should not close streams
- [SOLR-8453](#) - Solr should attempt to consume the request inputStream on errors as we cannot count on the container to do it
- [SOLR-8575](#) - Fix HDFSLogReader replay status numbers and a performance bug where we can reopen FSDataInputStream too often
- [SOLR-8578](#) - Successful or not, requests are not always fully consumed by Solrj clients and we count on HttpClient or the JVM
- [SOLR-8615](#) - Just like creating cores, we should use multiple threads when closing cores
- [SOLR-8633](#) - DistributedUpdateProcess processCommit/deleteByQuery calls finish on DUP and SolrCmdDistributor, which violates the lifecycle and can cause bugs
- [SOLR-8720](#) - ZkController#publishAndWaitForDownStates should use #publishNodeAsDown
- [SOLR-8771](#) - Multi-threaded core shutdown creates executor per core
- [SOLR-8855](#) - The HDFS BlockDirectory should not clean up its cache on shutdown
- [SOLR-8856](#) - Do not cache merge or 'read once' contexts in the hdfs block cache
- [SOLR-8857](#) - HdfsUpdateLog does not use configured or new default number of version buckets and is hard coded to 256
- [SOLR-8869](#) - Optionally disable printing field cache entries in SolrFieldCacheMBean
- [SPARK-10859](#) - [SQL] Fix stats of StringType in columnar cache
- [SPARK-10914](#) - UnsafeRow serialization breaks when two machines have different Oups size
- [SPARK-11009](#) - [SQL] Fix wrong result of Window function in cluster mode
- [SPARK-11537](#) - [SQL] Fix negative hours/minutes/seconds
- [SPARK-11737](#) - [SQL] Fix serialization of UTF8String with Kyro
- [SPARK-12617](#) - [PYSOLR] Move Py4jCallbackConnectionCleaner to Streaming, clean up the leak sockets of Py4J

- [SPARK-14477](#) - [BUILD] Allow custom mirrors for downloading artifacts in build/mvn
- [SQOOP-2847](#) - Sqoop --incremental + missing parent --target-dir reports success with no data

Issues Fixed in CDH 5.5.2

Known Issues Fixed

The following topics describe known issues fixed in CDH 5.5.2.

Apache Spark

Spark SQL cannot retrieve data from a partitioned Hive table

When reading from a partitioned Hive table, Spark SQL is not able to identify the column delimiter used, and reads the full record as the first column entry.

Cloudera Bug: CDH-37189

Workaround: None.

When using Spark on YARN, the driver reports misleading error messages

The Spark driver reports misleading error messages such as:

```
ERROR ErrorMonitor: AssociationError [akka.tcp://sparkDriver@...] ->
[akka.tcp://sparkExecutor@...]: Error [Association failed with [akka.tcp://sparkE
xecutor@...]]
[akka.remote.EndpointAssociationException: Association failed with
[akka.tcp://sparkExecutor@...]]
```

Cloudera Bug: CDH-33949

Workaround: Add the following property to the Spark log4j configuration file:

```
log4j.logger.org.apache.spark.rpc.akka.ErrorMonitor=FATAL..
```

Spark does not support rolling upgrades

Spark does not support rolling upgrades. Submitted Spark jobs may fail during upgrade. Jobs requiring new configuration properties will fail.

Cloudera Bug: CDH-33655

Workaround: Finish the upgrade, and then relaunch the Spark jobs.

Hue

Cannot query the `customers` table in Hue

To query the `customers` table, you must re-create the Parquet data for compatibility.

Bug: [HUE-3040](#)

Cloudera Bug: CDH-33974

Workaround: Update the parquet file of the `customers` table (`/user/hive/warehouse/customers/customers`) with the one attached to [HUE-3040](#).

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.5.2:

- [AVRO-1781](#) - Remove LogicalTypes cache
- [HADOOP-7713](#) - `dfs -count -q` should label output column
- [HADOOP-10406](#) - `TestIPC.testIpcWithReaderQueuing` may fail
- [HADOOP-10668](#) - Addendum patch to fix `TestZKFailoverController`
- [HADOOP-10668](#) - `TestZKFailoverControllerStress#testExpireBackAndForth` occasionally fails
- [HADOOP-11171](#) - Enable using a proxy server to connect to S3a
- [HADOOP-11218](#) - Add TLS 1.1, TLS 1.2 to KMS, HttpFS, SSLFactory
- [HADOOP-12269](#) - Update `aws-sdk` dependency to 1.10.6

- [HADOOP-12417](#) - TestWebDelegationToken failing with port in use
- [HADOOP-12418](#) - TestRPC.testRPCInterruptedSimple fails intermittently
- [HADOOP-12464](#) - Interrupted client may try to fail-over and retry
- [HADOOP-12468](#) - Partial group resolution failure should not result in user lockout
- [HADOOP-12474](#) - MiniKMS should use random ports for Jetty server by default
- [HADOOP-12568](#) - Update core-default.xml to describe posixGroups support
- [HADOOP-12573](#) - TestRPC.testClientBackOff failing
- [HADOOP-12584](#) - Disable browsing the static directory in HttpServer2
- [HADOOP-12584](#) - Revert - Disable browsing the static directory in HttpServer2
- [HADOOP-12584](#) - Disable browsing the static directory in HttpServer2
- [HADOOP-12604](#) - Exception may be swallowed in KMSClientProvider
- [HADOOP-12625](#) - Add a config to disable the /logs endpoints
- [HDFS-6101](#) - TestReplaceDatanodeOnFailure fails occasionally
- [HDFS-6533](#) - TestBPOfferService#testBasicFunctionalitytest fails intermittently
- [HDFS-6694](#) - TestPipelinesFailover.testPipelineRecoveryStress tests fail intermittently with various symptoms - debugging patch
- [HDFS-7553](#) - Fix the TestDFSUpgradeWithHA due to BindException
- [HDFS-7798](#) - Checkpointing failure caused by shared KerberosAuthenticator
- [HDFS-8647](#) - Abstract BlockManager's rack policy into BlockPlacementPolicy
- [HDFS-8722](#) - Optimize DataNode writes for small writes and flushes
- [HDFS-8772](#) - Fix TestStandbyIsHot#testDatanodeRestarts which occasionally fails
- [HDFS-8805](#) - Archival Storage: getStoragePolicy should not need superuser privilege
- [HDFS-9083](#) - Replication violates block placement policy
- [HDFS-9123](#) - Copying from the root to a subdirectory should be forbidden
- [HDFS-9160](#) - [OIV-Doc] : Missing details of 'delimited' for processor options
- [HDFS-9220](#) - Reading small file (< 512 bytes) that is open for append fails due to incorrect checksum
- [HDFS-9249](#) - NPE is thrown if an IOException is thrown in NameNode constructor
- [HDFS-9250](#) - Add precondition check to LocatedBlock#addCachedLoc
- [HDFS-9268](#) - fuse_dfs chown crashes when uid is passed as -1
- [HDFS-9273](#) - ACLs on root directory may be lost after NameNode restart
- [HDFS-9286](#) - HttpFs does not parse ACL syntax correctly for operation REMOVEACLENTRIES
- [HDFS-9295](#) - Add a thorough test of the full KMS code path
- [HDFS-9313](#) - Possible NullPointerException in BlockManager if no excess replica can be chosen
- [HDFS-9332](#) - Fix Precondition failures from NameNodeEditLogRoller while saving namespace
- [HDFS-9339](#) - Extend full test of KMS ACLs
- [HDFS-9364](#) - Unnecessary DNS resolution attempts when creating NameNodeProxies
- [HDFS-9410](#) - Some tests should always reset sysout and syserr
- [HDFS-9429](#) - Tests in TestDFSAdminWithHA intermittently fail with EOFException
- [HDFS-9438](#) - Only collect HDFS-6694 debug data on Linux, Mac, and Solaris
- [HDFS-9445](#) - DataNode may deadlock while handling a bad volume
- [HDFS-9470](#) - Encryption zone on root not loaded from fsimage after NameNode restart
- [HDFS-9474](#) - TestPipelinesFailover should not fail when printing debug message
- [MAPREDUCE-6191](#) - Improve clearing stale state of Java serialization testcase
- [MAPREDUCE-6233](#) - org.apache.hadoop.mapreduce.TestLargeSort.testLargeSort failed in trunk
- [MAPREDUCE-6549](#) - Multibyte delimiters with LineRecordReader cause duplicate records
- [MAPREDUCE-6550](#) - archive-logs tool changes log ownership to the Yarn user when using DefaultContainerExecutor
- [YARN-3564](#) - Fix TestContainerAllocation.testAMContainerAllocationWhenDNSUnavailable fails randomly
- [YARN-3768](#) - ArrayIndexOutOfBoundsException with empty environment variables
- [YARN-4235](#) - FairScheduler PrimaryGroup does not handle empty groups returned for a user
- [YARN-4310](#) - FairScheduler: Log skipping reservation messages at DEBUG level

- [YARN-4347](#) - Resource manager fails with Null pointer exception
- [YARN-4408](#) - Fix issue that NodeManager still reports negative running containers
- [HBASE-6617](#) - ReplicationSourceManager should be able to track multiple WAL paths
- [HBASE-12961](#) - Fix negative values in read and write region server metrics
- [HBASE-13134](#) - mutateRow and checkAndMutate apis don't throw region level exceptions
- [HBASE-13703](#) - ReplicateContext should not be a member of ReplicationSource
- [HBASE-13746](#) - list_replicated_tables command is not listing table in HBase shell
- [HBASE-13833](#) - LoadIncrementalHFile.doBulkLoad(Path, HTable) does not handle unmanaged connections when using SecureBulkLoad
- [HBASE-14003](#) - Work around JDK-8044053
- [HBASE-14205](#) - RegionCoprocessorHost System.nanoTime() performance bottleneck
- [HBASE-14283](#) - Reverse scan doesn't work with HFile inline index/bloom blocks
- [HBASE-14501](#) - NPE in replication with TDE
- [HBASE-14533](#) - Connection Idle time 1 second is too short and the connection is closed too quickly by the ChoreService. Increase it to the default (10 minutes) for testAll(). The patch is not committed upstream yet.
- [HBASE-14541](#) - TestHFileOutputFormat.testMRIncrementalLoadWithSplit failed due to too many splits and few retries
- [HBASE-14547](#) - Add more debug/trace to zk-procedure
- [HBASE-14621](#) - ReplicationLogCleaner stuck on RS crash
- [HBASE-14731](#) - Add -DuseMob option to ITBLL
- [HBASE-14809](#) - Grant / revoke namespace admin permission to group
- [HBASE-14923](#) - VerifyReplication should not mask the exception during result comparison
- [HBASE-14926](#) - Hung ThriftServer; no timeout on read from client; if client crashes, worker thread gets stuck reading
- [HBASE-15031](#) - Fix merge of MVCC and SequenceID performance regression in branch-1.0
- [HBASE-15032](#) - HBase shell scan filter string assumes UTF-8 encoding
- [HBASE-15035](#) - Bulkloading HFiles with tags that require splits do not preserve tags
- [HBASE-15104](#) - Occasional failures due to NotServingRegionException in IT tests
- [HIVE-7575](#) - GetTables thrift call is very slow
- [HIVE-7653](#) - Hive AvroSerDe does not support circular references in Schema
- [HIVE-9507](#) - Make "LATERAL VIEW inline(expression) mytable" tolerant to nulls
- [HIVE-10027](#) - Use descriptions from Avro schema files in column comments
- [HIVE-10048](#) - JDBC - Support SSL encryption regardless of Authentication mechanism
- [HIVE-10083](#) - SMBJoin fails in case one table is uninitialized
- [HIVE-10265](#) - Hive CLI crashes on != inequality
- [HIVE-10514](#) - Fix MiniCliDriver tests failure
- [HIVE-10687](#) - AvroDeserializer fails to deserialize evolved union fields
- [HIVE-10697](#) - ObjectInspectorConvertors#UnionConvertor does a faulty conversion
- [HIVE-11149](#) - Fix issue with sometimes HashMap in PerfLogger.java hangs
- [HIVE-11288](#) - Backport:Avro SerDe InstanceCache returns incorrect schema
- [HIVE-11513](#) - AvroLazyObjectInspector could handle empty data better
- [HIVE-11616](#) - DelegationTokenSecretManager reuses the same objectstore, which has concurrency issues
- [HIVE-11785](#) - Revert - Support escaping carriage return and new line for LazySimpleSerDe
- [HIVE-11785](#) - Support escaping carriage return and new line for LazySimpleSerDe
- [HIVE-11826](#) - 'hadoop.proxyuser.hive.groups' configuration does not prevent unauthorized user to access metastore
- [HIVE-11977](#) - Hive should handle an external Avro table with zero length files present
- [HIVE-12008](#) - Hive queries failing when using count(*) on column in view
- [HIVE-12058](#) - Change Hive script to record errors when calling hbase fails
- [HIVE-12188](#) - DoAs does not work properly in non-Kerberos secured HiveServer2
- [HIVE-12189](#) - The list in pushdownPreds of ppd.ExprWalkerInfo should not be allowed to grow very large

- [HIVE-12218](#) - Unable to create a like table for an HBase-backed table
- [HIVE-12250](#) - Zookeeper connection leaks in Hive's HBaseHandler
- [HIVE-12265](#) - Generate lineage info only if requested
- [HIVE-12268](#) - Context leaks deleteOnExit paths
- [HIVE-12278](#) - Skip logging lineage for explain queries
- [HIVE-12287](#) - Lineage for lateral view shows wrong dependencies
- [HIVE-12330](#) - Fix precommit Spark test part2
- [HIVE-12365](#) - Added resource path is sent to cluster as an empty string when externally removed
- [HIVE-12378](#) - Exception on HBaseSerDe.serialize binary field
- [HIVE-12388](#) - GetTables cannot get external tables when TABLE type argument is given
- [HIVE-12406](#) - HIVE-9500 introduced incompatible change to LazySimpleSerDe public interface
- [HIVE-12418](#) - HiveHBaseTableInputFormat.getRecordReader() causes Zookeeper connection leak
- [HIVE-12505](#) - Backport: Insert overwrite in same encrypted zone silently fails to remove some existing files
- [HIVE-12566](#) - Incorrect result returns when using COALESCE in WHERE condition with LEFT JOIN
- [HIVE-12713](#) - Miscellaneous improvements in driver compile and execute logging
- [HIVE-12784](#) - Group by SemanticException: Invalid column reference
- [HIVE-12788](#) - Setting hive.optimize.union.remove to TRUE will break UNION ALL with aggregate functions
- [HIVE-12795](#) - Vectorized execution causes ClassCastException
- [HUE-2664](#) - Revert - [jobbrowser] Fix fetching logs from job history server
- [HUE-2997](#) - [oozie] Easier usage of email action when workflow fails
- [HUE-3035](#) - [beeswax] Optimize sample data query for partitioned tables
- [HUE-3036](#) - [beeswax] Revert get_tables to use Thrift API GetTables
- [HUE-3091](#) - [oozie] Do not remove extra new lines from email action body
- [IMPALA-1459](#) - Fix migration/assignment of On-clause predicates inside inline views.
- [IMPALA-2103](#) - Fix flaky test_impersonation test
- [IMPALA-2113](#) - Handle error when distinct and aggregates are used with a having clause
- [IMPALA-2225](#) - Handle error when star based select item and aggregate are incorrectly used
- [IMPALA-2226](#) - Throw AnalysisError if table properties are too large
- [IMPALA-2273](#) - Make MAX_PAGE_HEADER_SIZE configurable
- [IMPALA-2473](#) - Reduce scanner memory usage
- [IMPALA-2535](#) - PAGG hits mem_limit when switching to I/O buffers
- [IMPALA-2558](#) - DCHECK in Parquet scanner after block read error
- [IMPALA-2559](#) - Fix check failed: sorter_runs_.back()->is_pinned_
- [IMPALA-2591](#) - DataStreamSender::Send() does not return an error status if SendBatch() failed
- [IMPALA-2598](#) - Re-enable SSL and Kerberos on server-server
- [IMPALA-2612](#) - Free local allocations once for every row batch when building hash tables
- [IMPALA-2614](#) - Don't ignore Status returned by DataStreamRecvr::CreateMerger()
- [IMPALA-2624](#) - Increase fs.trash.interval to 24 hours for test suite
- [IMPALA-2630](#) - Skip TestParquet.test_continue_on_error when using old aggs/joins
- [IMPALA-2643](#) - Prevent migrating incorrectly inferred identity predicates into inline views
- [IMPALA-2648](#) - Avoid sending large partition stats objects over thrift
- [IMPALA-2695](#) - Fix GRANTs on URIs with uppercase letters
- [IMPALA-2722](#) - Free local allocations per row batch in non-partitioned AGG and HJ
- [IMPALA-2731](#) - Refactor MemPool usage in HBase scan node
- [IMPALA-2747](#) - Thrift-client cleans openSSL state before using it in the case of the catalog
- [IMPALA-2776](#) - Remove escapechartesttable and associated tests
- [IMPALA-2812](#) - Remove additional test referencing escapecharstesttable
- [IMPALA-2829](#) - SEGV in AnalyticEvalNode touching NULL input_stream_
- [KITE-1089](#) - ReadAvroContainer morphline command should work even if the Avro writer schema of each input file is different

- [KITE-1097](#) - Add method to read the name of a Morphline command
- [OOZIE-2030](#) - Configuration properties from global section is not getting set in Hadoop job conf when using sub-workflow action in Oozie workflow.xml
- [OOZIE-2365](#) - Oozie fails to start when SMTP password not set
- [OOZIE-2380](#) - Oozie Hive action failed with wrong tmp path
- [OOZIE-2397](#) - LAST_ONLY and NONE don't properly handle READY actions
- [OOZIE-2413](#) - Kerberos credentials can expire if the KDC is slow to respond
- [OOZIE-2439](#) - FS Action no longer uses name-node from global section or default NN
- [OOZIE-2441](#) - SubWorkflow action with propagate-configuration but no global section throws NPE on submit
- [PIG-3641](#) - Split "otherwise" producing incorrect output when combined with ColumnPruning
- [SENTRY-565](#) - Improve performance of filtering Hive SHOW commands
- [SENTRY-835](#) - Drop table leaves a connection open when using metastorelistener
- [SENTRY-902](#) - SimpleDBProviderBackend should retry the authorization process properly
- [SENTRY-936](#) - getGroup and getUser should always return original HDFS values for paths in prefix which are not managed by Sentry
- [SENTRY-944](#) - Setting HDFS rules on Sentry-managed HDFS paths should not affect original HDFS rules
- [SENTRY-953](#) - External Partitions which are referenced by more than one table can cause some unexpected behavior with Sentry HDFS sync
- [SENTRY-957](#) - Exceptions in MetastoreCacheInitializer should probably not prevent HMS from starting up
- [SENTRY-960](#) - Blacklist reflect, java_method using hive.server2.builtin.udf.blacklist
- [SENTRY-988](#) - Let SentryAuthorization setter path always fall through and update HDFS
- [SENTRY-994](#) - SentryAuthorizationInfoX should override isSentryManaged
- [SOLR-6443](#) - backportDisable test that fails on Jenkins until we can determine the problem
- [SOLR-7049](#) - LIST Collections API call should be processed directly by the CollectionsHandler instead of the OverseerCollectionProcessor
- [SOLR-7989](#) - After a new leader is elected it, it should ensure it's state is ACTIVE if it has already registered with ZooKeeper
- [SOLR-8075](#) - Fix faulty implementation
- [SOLR-8152](#) - Overseer Task Processor/Queue can miss responses, leading to timeouts
- [SOLR-8223](#) - Avoid accidentally swallowing OutOfMemoryError
- [SOLR-8288](#) - DistributedUpdateProcessor#doFinish should explicitly check and ensure it does not try to put itself into LIR
- [SOLR-8353](#) - Support regex for skipping license checksums
- [SOLR-8367](#) - Fix the LeaderInitiatedRecovery 'all replicas participate' fail-safe
- [SOLR-8372](#) - backportCanceled recovery can lead to data loss
- [SOLR-8535](#) - Support forcing define-lucene-javadoc-url to be local
- [SPARK-5569](#) - [STREAMING] Fix ObjectInputStreamWithLoader for supporting load array classes
- [SPARK-8029](#) - Robust shuffle writer
- [SPARK-9735](#) - [SQL] Respect the user specified schema than the infer partition schema for HadoopFsRelation
- [SPARK-10648](#) - Oracle dialect to handle nonspecific numeric types
- [SPARK-10865](#) - [SPARK-10866] [SQL] Fix bug of ceil/floor, which should returns long instead of the Double type
- [SPARK-11105](#) - [YARN] Distribute log4j.properties to executors
- [SPARK-11126](#) - [SQL] Fix the potential flaky test
- [SPARK-11126](#) - [SQL] Fix a memory leak in SQLListener._stageIdToStageMetrics
- [SPARK-11246](#) - [SQL] Table cache for Parquet broken in 1.5
- [SPARK-11453](#) - [SQL] Append data to partitioned table will messes up the result
- [SPARK-11484](#) - [WEBUI] Using proxyBase set by Spark AM
- [SPARK-11786](#) - [CORE] Tone down messages from akka error monitor
- [SPARK-11799](#) - [CORE] Make it explicit in executor logs that uncaught exceptions are thrown during executor shutdown
- [SPARK-11929](#) - [CORE] Make the repl log4j configuration override the root logger

- [SQOOP-2745](#) - Using datetime column as a splitter for Oracle no longer works
- [SQOOP-2767](#) - Test is failing SystemImportTest
- [SQOOP-2783](#) - Query import with parquet fails on incompatible schema
- [SQOOP-2422](#) - Sqoop2: Test TestJSONIntermediateDataFormat is failing on JDK8

Issues Fixed in CDH 5.5.1

The following issues have been fixed in CDH 5.5.1:

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera (“Cloudera Products”), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled “Arbitrary remote code execution with InvokerTransformer” and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Apache HBase

Data may not be replicated to slave cluster if multiway multiplicity is set to greater than 1.

Bug: [HBASE-13703](#), [HBASE-6617](#), [HBASE-14501](#), CDH-30208.

Issues Fixed in CDH 5.5.0

The following topics describe known issues fixed in CDH 5.5.0.

Apache Flume

Fix FD leak in AsyncHBaseSink.

Bug: [FLUME-2738](#)

Fix for Kerberos configuration error when using short names.

Bug: [FLUME-2749](#)

Fix NullPointerException in KafkaSourceCounter.

Bug: [FLUME-2672](#)

Fix for Tail Directory Source FileNotFoundException.

Bug: [FLUME-2773](#)

Fix for Kafka Channel timeout property handling.

Bug: [FLUME-2734](#)

Apache Hadoop
YARN/MapReduce

Incorrect headroom leads to deadlock between mappers and reducers.

Bug: [MAPREDUCE-6302](#)

Blacklisting Support for Scheduling ApplicationMasters

When an ApplicationMaster fails, and the NodeManager on the same host has not yet been blacklisted, the framework should route the second ApplicationMaster attempt to a NodeManager on a different host.

Bug: [YARN-2005](#)

Cloudera Bug: CDH-25492

Apache HBase

HBase mlock agent is not included as part of CDH 5.x releases.

Starting in CDH 5.0, the native mlock daemons were not included in CDH HBase. CDH 5.5 restores the daemon in both parcels and packages.

Cloudera Bug: CDH-30196, CDH-30208

Apache Hive

Parquet Predicate pushdown for float types does not work

The Parquet predicate builder should use PrimitiveTypeName type to construct a predicate leaf instead of the type provided by PredicateLeaf.

Bug: [HIVE-11504](#)

Cloudera Bug: CDH-30495

LineageCtx should release all resources at clear()

Some maps are not released with the clear() method and can cause a memory leak.

Bug: [HIVE-12225](#)

Cloudera Bug: CDH-33265

LEFT JOIN query plan outputs wrong column when using subquery

Incorrect results may arise if a LEFT OUTER JOIN is combined with a subquery.

Bug: [HIVE-9613](#)

Cloudera Bug: CDH-28587

Map-side aggregation is extremely slow

Map-side aggregation on columns with double type is extremely slow due to [HIVE-7041](#).

Bug: [HIVE-11502](#)

Cloudera Bug: CDH-30387

Lineage does not work with dynamic partitioning query

An error message displays after running a dynamic partitioning query: ERROR : Result schema has 2 fields, but we don't get as many dependencies.

Bug: [HIVE-11834](#)

Cloudera Bug: CDH-31566

DROP PARTITION in encrypted zone does not remove data from HDFS

An ALTER TABLE query to DROP PARTITION removes the partition metadata from HDFS but not the data.

Bug: [HIVE-10910](#)

Cloudera Bug: CDH-29014*DROP TABLE with qualified table name ignores database name when checking partitions*

DDLTask.dropTable() uses an older version of Hive.getPartitionNames(), which takes in a single string for the table name, instead of the database and table names.

Bug: [HIVE-10421](#)**Cloudera Bug:** CDH-31183*INSERT INTO statement may expose data that should be encrypted*

INSERT INTO <table> VALUES() uses a temporary table; the data in temporary tables is stored under hive.exec.scratchdir which is not usually encrypted.

Bug: [HIVE-10658](#)**Cloudera Bug:** CDH-29235*Whitelist restrictions do not get initialized in new copy of HiveConf*

Whitelist restrictions use a regex pattern in HiveConf, but when a new HiveConf object copy is created, the regex pattern is not initialized in the new HiveConf copy.

Bug: [HIVE-10465](#)**Cloudera Bug:** CDH-31184*RuntimeException when vectorization is enabled with binary data*

A RuntimeException is thrown when vectorization is enabled and binary data is in the GROUP BY clause.

Bug: [HIVE-9908](#)**Cloudera Bug:** CDH-31182*Hive may return wrong results in some queries with PTF function*

The select statement has an extra column with a PTF operator that is skewing results.

Bug: [HIVE-11604](#)**Cloudera Bug:** CDH-30669*HiveServer2 leaks Hive Metastore Connections*

HiveServer2 uses threadlocal to cache Hive Metastore (HMS) Thrift client in class Hive. When the thread dies, the HMS client does not close. So the connection to the HMS client leaks.

Bug: [HIVE-10956](#)**Cloudera Bug:** CDH-28441*Remote Spark Client has a memory leak*

In Remote Spark Client (RSC), MapWork/ReduceWork tasks build up until an OutOfMemoryException is thrown.

Bug: [HIVE-10006](#)**Cloudera Bug:** CDH-26327*Replication factor is not properly set in SparkHashTableSinkOperator*

The default replication factors (3) affects the Map Join performance of small files.

Bug: [HIVE-11109](#)**Cloudera Bug:** CDH-28903

Hive LDAP Authenticator should allow users to set Domain without the base Distinguished Name

When the base distinguished name (DN) is not configured but only the Domain has been set in hive-site.xml, the LDAP authentication provider cannot locate the user in the directory. Authentication fails in such cases.

Bug: [HIVE-12007](#)

Cloudera Bug: CDH-23561

Hive should support additional LDAP authentication parameters

Currently, Hive only has the following authenticator parameters for LDAP authentication for HiveServer2:

```
<property>
  <name>hive.server2.authentication</name>
  <value>LDAP</value>
</property>
<property>
  <name>hive.server2.authentication.ldap.url</name>
  <value>ldap://our_ldap_address</value>
</property>
```

Other LDAP properties need to be included as part of Hive-LDAP authentication, for example:

```
group search base -> dc=domain,dc=com
group search filter -> member={0}
user search base -> dc=domain,dc=com
user search filter -> sAMAccountName={0}
list of valid user groups -> group1,group2,group3
```

Bug: [HIVE-7193](#)

Cloudera Bug: CDH-31454

Aggregate functions used as window functions can fail in various ways

- [HIVE-11817](#), CDH-28903: Window function max() fails with NullPointerException.
- [HIVE-10702](#), CDH-31128: COUNT(*) over windowing x preceding and y preceding returns unexpected results.
- [HIVE-10826](#), CDH-31958: Support min()/max() functions over x preceding and y preceding windowing.

Apache Spark

Attempts to access secure HBase from Spark executors fail when authenticating to the metastore.

An exception like the following occurs when you attempt to access kerberized HBase instance from a Spark executor.

```
GSSEException: No valid credentials provided
(Mechanism level: Failed to find any Kerberos tgt)
```

The root cause is that the HBase Kerberos authentication token is not sent to the Spark executor.

Bug: [SPARK-6918](#)

Workaround: None.

The shuffle service fails on NodeManager restarts and kills all running Spark applications

In CDH 5.4.0 through CDH 5.4.4, the shuffle service is on by default. Because it fails in NodeManager restarts, in CDH 5.4.5, and higher, the shuffle service is off by default. Dynamic allocation requires that the shuffle service be turned on.

Bug: [SPARK-9439](#)

Cloudera Bug: CDH-29772

Workaround: In CDH 5.4.5 and higher, enable the shuffle service when using dynamic allocation.

Spark not automatically picking up hive-site.xml

When you run Spark on YARN, the client `hive-site.xml` does not get picked up automatically by `spark-submit`.

Bug: [SPARK-2669](#)

Cloudera Bug: CDH-23780

Workaround: Do one of the following, depending on which deployment mode you are running in:

- Client - set `HADOOP_CONF_DIR` to `/etc/hive/conf/` (or the directory where `hive-site.xml` is located).
- Cluster - add `--files=/etc/hive/conf/hive-site.xml` (or the path for `hive-site.xml`) to the `spark-submit` script.

*Apache Sentry (incubating)**Synchronize calls in SentryClient and create Sentry client once per request in SimpleDBProvider*

Adds proper locking to the SentryClient and reduces the number of SentryClients created within a single request in the SimpleDbProvider (used by Hive). This fixes issues that may have caused transient permission failures and out of memory conditions.

Bug: [SENTRY-893](#)

Sentry-HDFS sync events should treat database and table names as case-insensitive

Sentry-HDFS Sync was treating database and table names as case-sensitive. This led to incorrect or missing ACLs being applied as part of the sync operation if the DDL operations used a different case for the catalog objects.

Bug: [SENTRY-885](#)

Hive drop database operation removes the Sentry privileges, even if drop operation fails

Even if the Hive drop database operation fails, the Sentry privileges on that database will be removed.

Bug: [SENTRY-669](#)

Nested queries in Hive on views incorrectly enforce base table privileges instead of view privileges

Nested queries in Hive on views incorrectly enforce base table privileges instead of view privileges. This leads to to Hive query failures due to insufficient privileges.

Bug: [SENTRY-619](#)

*Apache ZooKeeper**BinaryInputArchive readString should check length before allocating memory*

This fixed a possible OutOfMemoryError when malformed packets were sent to the ZooKeeper server.

Bug: [ZOOKEEPER-2146](#)

Cloudera Bug: CDH-25733

Workaround: Upgrade to CDH 5.5.

*Cloudera Search**The GoLive Function Does not Support Running As a Configurable User*

After using `--go-live` mode with the `MapReduceIndexerTool` and `HBaseMapReduceIndexerTool`, depending on group mappings and the configured HDFS umask, Solr may not have been able to read the results of the indexing job.

With Search for CDH 5.5 and later, the `MapReduceIndexerTool` and `HBaseMapReduceIndexerTool` includes updated `--go-live` functionality. The indexers now automatically update HDFS ACLs for the specified output directory, giving Solr permission to read the indexer results.

See [MapReduceIndexerTool](#) and [HBaseMapReduceIndexerTool](#) for more information.

Bug: None.

Workaround: Do not use the `--go-live` mode with `MapReduceIndexerTool` and `HBaseMapReduceIndexerTool` or use a less restrictive umask.

MapReduceIndexerTool fails to Index Documents When Sentry Is Enabled

Prior to CDH 5.5, when Sentry was enabled, the MapReduceIndexerTool was unable to index data even if the user was authorized to write to the collection according to Sentry permissions. This limitation occurred because, by default, the MapReduceIndexerTool used the underlying collection's `solrconfig.xml` from ZooKeeper to build the index using its EmbeddedSolrServers. But the embedded servers are not properly configured to use Sentry, so this process failed.

With Search for CDH 5.5, the MapReduceIndexerTool uses a default `solrconfig.xml` that is appropriate for the vast majority of collection configurations. With this configuration, the MapReduceIndexerTool is able to index data, even if Sentry is enabled. Note that this default configuration does not include any `updateRequestProcessorChains`; if your configuration requires an `updateRequestProcessorChain`, you can tell the MapReduceIndexerTool to use the configuration from ZooKeeper by specifying `--use-zk-solrconfig.xml` or from local disk by specifying `--solr-home-dir`.

Bug: None.

Workaround: To address this issue, configure the MapReduceIndexerTool to run without Sentry restrictions. This does not compromise security because this only affects the "embedded" Solr Servers in the job that are used to build the offline index; Solr's Sentry permissions are still checked when the data is merged into the cluster via `--go-live`.

Here are two ways to enable indexing:

1. If your environment uses the default configuration files, use `solrconfig.xml` for indexing jobs, rather than `solrconfig.xml.secure`. Use the `--solr-home-dir` option to specify the directory containing `solrconfig.xml`, causing the job to run with Sentry disabled.
2. Alternately, you can comment out the following line:

```
<str name="update.chain">updateIndexAuthorization</str>
```

This line must be commented out and the change saved in the `solrconfig` file used by the machine running the indexing job.

Issues Fixed in CDH 5.4.x

The following topics describe known issues fixed in CDH 5.4.x, from newest to oldest release.

Issues Fixed in CDH 5.4.11

CDH 5.4.11 fixes the following issues.

- [FLUME-2891](#) - Revert FLUME-2712 and FLUME-2886
- [FLUME-2908](#) - NetcatSource - SocketChannel not closed when session is broken
- [HADOOP-8436](#) - NPE In getLocalPathForWrite (path, conf) when the required context item is not configured
- [HADOOP-8437](#) - getLocalPathForWrite should throw IOException for invalid paths
- [HADOOP-8751](#) - NPE in Token.toString() when Token is constructed using null identifier
- [HADOOP-8934](#) - Shell command ls should include sort options
- [HADOOP-10048](#) - LocalDirAllocator should avoid holding locks while accessing the filesystem
- [HADOOP-10971](#) - Add -C flag to make `hadoop fs -ls` print filenames only
- [HADOOP-11901](#) - BytesWritable fails to support 2G chunks due to integer overflow
- [HADOOP-12252](#) - LocalDirAllocator should not throw NPE with empty string configuration
- [HADOOP-12269](#) - Update aws-sdk dependency to 1.10.6
- [HADOOP-12787](#) - KMS SPNEGO sequence does not work with WEBHDFS
- [HADOOP-12841](#) - Update s3-related properties in core-default.xml.
- [HADOOP-12901](#) - Add warning log when KMSClientProvider cannot create a connection to the KMS server.
- [HADOOP-12972](#) - Lz4Compressor#getLibraryName returns the wrong version number
- [HADOOP-13079](#) - Add -q option to ls to print ? instead of non-printable characters
- [HADOOP-13132](#) - Handle ClassCastException on AuthenticationException in LoadBalancingKMSClientProvider
- [HADOOP-13155](#) - Implement TokenRenewer to renew and cancel delegation tokens in KMS
- [HADOOP-13251](#) - Authenticate with Kerberos credentials when renewing KMS delegation token

- [HADOOP-13255](#) - KMSSClientProvider should check and renew tgt when doing delegation token operations
- [HADOOP-13263](#) - Reload cached groups in background after expiry.
- [HADOOP-13457](#) - Remove hardcoded absolute path for shell executable.
- [HDFS-4660](#) - Block corruption can happen during pipeline recovery
- [HDFS-8211](#) - DataNode UUID is always null in the JMX counter.
- [HDFS-8451](#) - DFSClient probe for encryption testing interprets empty URI property for enabled
- [HDFS-8496](#) - Calling stopWriter() with FSDataSetImpl lock held may block other threads
- [HDFS-8576](#) - Lease recovery should return true if the lease can be released and the file can be closed
- [HDFS-8722](#) - Optimize datanode writes for small writes and flushes
- [HDFS-9085](#) - Show renewer information in DelegationTokenIdentifier#toString
- [HDFS-9220](#) - Reading small file (< 512 bytes) that is open for append fails due to incorrect checksum
- [HDFS-9276](#) - Failed to Update HDFS Delegation Token for long running application in HA mode
- [HDFS-9466](#) - TestShortCircuitCache#testDataXceiverCleansUpSlotsOnFailure is flaky
- [HDFS-9589](#) - Block files which have been hardlinked should be duplicated before the DataNode appends to the them
- [HDFS-9700](#) - DFSClient and DFSOutputStream should set TCP_NODELAY on sockets for DataTransferProtocol
- [HDFS-9732](#) - Improve DelegationTokenIdentifier.toString() for better logging
- [HDFS-9939](#) - Increase DecompressorStream skip buffer size
- [HDFS-9949](#) - Add a test case to ensure that the DataNode does not regenerate its UUID when a storage directory is cleared
- [HDFS-10267](#) - Extra "synchronized" on FsDataSetImpl#recoverAppend and FsDataSetImpl#recoverClose
- [HDFS-10360](#) - DataNode might format directory and lose blocks if current/VERSION is missing.
- [HDFS-10381](#) - , DataStreamer DataNode exclusion log message should be warning.
- [MAPREDUCE-4785](#) - TestMRApp occasionally fails
- [MAPREDUCE-6580](#) - Test failure: TestMRJobsWithProfiler
- [YARN-2871](#) - TestRMRestart#testRMRestartGetApplicationList sometimes fails in trunk
- [YARN-3727](#) - Check if the directory exists before using it for localization
- [YARN-4168](#) - Fixed a failing test TestLogAggregationService.testLocalFileDeletionOnDiskFull
- [YARN-4354](#) - Public resource localization fails with NPE
- [YARN-4717](#) - TestResourceLocalizationService.testPublicResourceInitializesLocalDir fails Intermittently due to IllegalArgumentException from cleanup
- [YARN-5048](#) - DelegationTokenRenewer#skipTokenRenewal might throw NPE
- [HBASE-6617](#) - ReplicationSourceManager should be able to track multiple WAL paths (ADDENDUM)
- [HBASE-11625](#) - Verifies data before building HFileBlock. - Adds HFileBlock.Header class which contains information about location of fields. Testing: Adds CorruptedFSReaderImpl to TestChecksum.
- [HBASE-11927](#) - Use Native Hadoop Library for HFile checksum.
- [HBASE-14155](#) - StackOverflowError in reverse scan
- [HBASE-14359](#) - HTable#close will hang forever if unchecked error/exception thrown in AsyncProcess#sendMultiAction
- [HBASE-14730](#) - region server needs to log warnings when there are attributes configured for cells with hfile v2
- [HBASE-14759](#) - Avoid using Math.abs when selecting SyncRunner in FSHLog
- [HBASE-15234](#) - Don't abort ReplicationLogCleaner on ZooKeeper errors
- [HBASE-15456](#) - CreateTableProcedure/ModifyTableProcedure needs to fail when there is no family in table descriptor
- [HBASE-15479](#) - No more garbage or beware of autoboxing
- [HBASE-15582](#) - SnapshotManifestV1 too verbose when there are no regions
- [HBASE-15707](#) - ImportTSV bulk output does not support tags with hfile.format.version=3
- [HBASE-15746](#) - Remove extra RegionCoprocesor preClose() in RSRpcServices#closeRegion
- [HBASE-15811](#) - Batch Get after batch Put does not fetch all Cells We were not waiting on all executors in a batch to complete. The test for no-more-executors was damaged by the 0.99/0.98.4 fix "HBASE-11403 Fix race conditions around Object#notify"

- [HBASE-15925](#) - provide default values for hadoop compat module related properties that match default hadoop profile.
- [HBASE-16207](#) - can't restore snapshot without "Admin" permission
- [HIVE-9499](#) - hive.limit.query.max.table.partition makes queries fail on non-partitioned tables
- [HIVE-10048](#) - JDBC - Support SSL encryption regardless of Authentication mechanism
- [HIVE-10303](#) - HIVE-9471 broke forward compatibility of ORC files
- [HIVE-10685](#) - Alter table concatenate operator will cause duplicate data
- [HIVE-10925](#) - Non-static threadlocals in metastore code can potentially cause memory leak
- [HIVE-11031](#) - ORC concatenation of old files can fail while merging column statistics
- [HIVE-11054](#) - Handle varchar/char partition columns in vectorization
- [HIVE-11243](#) - Changing log level in Utilities.getBaseWork
- [HIVE-11408](#) - HiveServer2 is leaking ClassLoaders when add jar / temporary functions are used due to constructor caching in Hadoop ReflectionUtils
- [HIVE-11427](#) - Location of temporary table for CREATE TABLE SELECT broken by HIVE-7079.
- [HIVE-11488](#) - Combine the following jiras for "Support sessionId and queryId logging" Add sessionId and queryId info to HS2 log (Aihua Xu, reviewed by Szehon Ho) HIVE-12456: QueryId can't be stored in the configuration of the SessionState since multiple queries can run in a single session
- [HIVE-11583](#) - When PTF is used over a large partitions result could be corrupted
- [HIVE-11747](#) - Unnecessary error log is shown when executing a "INSERT OVERWRITE LOCAL DIRECTORY" cmd in the embedded mode
- [HIVE-11827](#) - STORED AS AVRO fails SELECT COUNT(*) when empty
- [HIVE-11919](#) - Hive Union Type Mismatch
- [HIVE-12354](#) - MapJoin with double keys is slow on MR
- [HIVE-12431](#) - Support timeout for compile lock
- [HIVE-12481](#) - Occasionally "Request is a replay" will be thrown from HS2
- [HIVE-12635](#) - Hive should return the latest hbase cell timestamp as the row timestamp value
- [HIVE-12958](#) - Make embedded Jetty server more configurable
- [HIVE-13200](#) - Aggregation functions returning empty rows on partitioned columns
- [HIVE-13251](#) - hive can't read the decimal in AVRO file generated from previous version
- [HIVE-13285](#) - Orc concatenation may drop old files from moving to final path
- [HIVE-13286](#) - Query ID is being reused across queries
- [HIVE-13462](#) - HiveResultSetMetaData.getPrecision() fails for NULL columns
- [HIVE-13527](#) - Using deprecated APIs in HBase client causes zookeeper connection leaks
- [HIVE-13570](#) - Some queries with Union all fail when CBO is off
- [HIVE-13932](#) - Hive SMB Map Join with small set of LIMIT failed with NPE
- [HIVE-13953](#) - Issues in HiveLockObject equals method
- [HIVE-13991](#) - Union All on view fail with no valid permission on underneath table
- [HIVE-14118](#) - Make the alter partition exception more meaningful
- [HUE-3185](#) - [oozie] Avoid extra API calls for parent information in workflow dashboard
- [HUE-3185](#) - Revert "[oozie] Avoid extra API calls for parent information in workflow dashboard"
- [HUE-3185](#) - [oozie] Avoid extra API calls for parent information in workflow dashboard
- [HUE-3437](#) - [core] PamBackend does not honor ignore_username_case
- [IMPALA-2378](#) - check proc mem limit before preparing fragment
- [IMPALA-2612](#) - Free local allocations once for every row batch when building hash tables.
- [IMPALA-2711](#) - Fix memory leak in Rand().
- [IMPALA-2722](#) - Free local allocations per row batch in non-partitioned AGG and HJ
- [OOZIE-2429](#) - TestEventGeneration test is unreliable
- [OOZIE-2466](#) - Repeated failure of TestMetricsInstrumentation.testSamplers
- [OOZIE-2486](#) - TestSLAEventsGetForFilterJPAExecutor is unreliable
- [SENTRY-780](#) - HDFS Plugin should not execute path callbacks for views

- [SENTRY-1184](#) - Clean up HMsPaths.renameAuthzObject
- [SENTRY-1292](#) - Reorder DBModelAction EnumSet
- [SENTRY-1293](#) - Avoid converting string permission to Privilege object
- [SOLR-6631](#) - DistributedQueue spinning on calling zookeeper getChildren()
- [SOLR-6820](#) - Make the number of version buckets used by the UpdateLog configurable as increasing beyond the default 256 has been shown to help with high volume indexing performance in SolrCloudIncrease the default number of buckets to 65536 instead of 256fix numVersionBuckets name attribute in configsets
- [SOLR-7332](#) - Initialize the highest value for all version buckets with the max value from the index or recent updates to avoid unnecessary lookups to the index to check for reordered updates when processing new documents.
- [SOLR-7587](#) - TestSpellCheckResponse stalled and never timed out
- [SOLR-7625](#) - Version bucket seed not updated after new index is installed on a replica
- [SOLR-8152](#) - Overseer Task Processor/Queue can miss responses, leading to timeouts
- [SOLR-8451](#) - Fix backport
- [SOLR-8451](#) - We should not call method.abort in HttpSolrClient or HttpSolrCall#remoteQuery and HttpSolrCall#remoteQuery should not close streams.
- [SOLR-8453](#) - Solr should attempt to consume the request inputStream on errors as we cannot count on the container to do it.
- [SOLR-8578](#) - Successful or not, requests are not always fully consumed by Solrj clients and we count on HttpClient or the JVM.
- [SOLR-8633](#) - DistributedUpdateProcess processCommit/deleteByQuery calls finish on DUP and SolrCmdDistributor, which violates the lifecycle and can cause bugs.
- [SOLR-8683](#) - Tune down stream closed logging
- [SOLR-8683](#) - Always consume the full request on the server, not just in the case of an error.
- [SOLR-8855](#) - The HDFS BlockDirectory should not clean up its cache on shutdown.
- [SOLR-8856](#) - Do not cache merge or 'read once' contexts in the hdfs block cache.
- [SOLR-8857](#) - HdfsUpdateLog does not use configured or new default number of version buckets and is hard coded to 256.
- [SOLR-8869](#) - Optionally disable printing field cache entries in SolrFieldCacheMBean
- [SPARK-12087](#) - Create new JobConf for every batch in saveAsHadoopFiles

Issues Fixed in CDH 5.4.10

CDH 5.4.10 fixes the following issues.

Apache Hadoop

FSImage may get corrupted after deleting snapshot

Bug: [HDFS-9406](#)

Cloudera Bug: CDH-33224

When deleting a snapshot that contains the last record of a given INode, the fsimage may become corrupt because the create list of the snapshot diff in the previous snapshot and the child list of the parent INodeDirectory are not cleaned.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.4.10:

- [FLUME-2712](#) - Optional channel errors slows down the Source to Main channel event rate
- [FLUME-2886](#) - Optional Channels can cause OOMs
- [HADOOP-10406](#) - TestIPC.testIpcWithReaderQueuing may fail
- [HADOOP-10668](#) - TestZKFailoverControllerStress#testExpireBackAndForth occasionally fails
- [HADOOP-11218](#) - Add TLSv1.1,TLSv1.2 to KMS, HttpFS, SSLFactory
- [HADOOP-12200](#) - TestCryptoStreamsWithOpensslAesCtrCryptoCodec should be skipped in non-native profile
- [HADOOP-12240](#) - Fix tests requiring native library to be skipped in non-native profile
- [HADOOP-12280](#) - Skip unit tests based on Maven profile rather than NativeCodeLoader.isNativeCodeLoaded

- [HADOOP-12417](#) - TestWebDelegationToken failing with port in use
- [HADOOP-12418](#) - TestRPC.testRPCInterruptedSimple fails intermittently
- [HADOOP-12464](#) - Interrupted client may try to fail-over and retry
- [HADOOP-12468](#) - Partial group resolution failure should not result in user lockout.
- [HADOOP-12474](#) - MiniKMS should use random ports for Jetty server by default
- [HADOOP-12559](#) - KMS connection failures should trigger TGT renewal
- [HADOOP-12604](#) - Exception may be swallowed in KMSClientProvider.
- [HADOOP-12605](#) - Fix intermittent failure of TestIPC.testIpcWithReaderQueuing
- [HADOOP-12668](#) - Support excluding weak Ciphers in HttpServer2 through ssl-server.conf
- [HADOOP-12682](#) - Fix TestKMS#testKMSRestart* failure
- [HADOOP-12699](#) - TestKMS#testKMSProvider intermittently fails during 'test rollover draining'
- [HADOOP-12715](#) - TestValueQueue#testgetAtMostPolicyALL fails intermittently
- [HADOOP-12736](#) - TestTimedOutTestsListener#testThreadDumpAndDeadlocks sometimes times out
- [HADOOP-12788](#) - OpensslAesCtrCryptoCodec should log which random number generator is used
- [HDFS-6533](#) - TestBPOfferService#testBasicFunctionalitytest fails intermittently
- [HDFS-7553](#) - fix the TestDFSUpgradeWithHA due to BindException
- [HDFS-8647](#) - Abstract BlockManager's rack policy into BlockPlacementPolicy
- [HDFS-9083](#) - Replication violates block placement policy
- [HDFS-9092](#) - Nfs silently drops overlapping write requests and causes data copying to fail
- [HDFS-9289](#) - Make DataStreamer#block thread safe and verify genStamp in commitBlock
- [HDFS-9313](#) - Possible NullPointerException in BlockManager if no excess replica can be chosen
- [HDFS-9347](#) - Invariant assumption in TestQuorumJournalManager.shutdown() is wrong
- [HDFS-9358](#) - TestNodeCount#testNodeCount timed out
- [HDFS-9406](#) - FSImage may get corrupted after deleting snapshot
- [HDFS-9445](#) - Datanode may deadlock while handling a bad volume
- [HDFS-9688](#) - Test the effect of nested encryption zones in HDFS downgrade
- [HDFS-9721](#) - Allow Delimited PB OIV tool to run upon fsimage that contains INodeReference
- [MAPREDUCE-6302](#) - Incorrect headroom can lead to a deadlock between map and reduce allocations
- [MAPREDUCE-6460](#) - TestRMContainerAllocator.testAttemptNotFoundCausesRMCommunicatorException fails
- [YARN-2902](#) - Killing a container that is localizing can orphan resources in the DOWNLOADING state
- [YARN-4155](#) - TestLogAggregationService.testLogAggregationServiceWithInterval failing
- [YARN-4204](#) - ConcurrentModificationException in FairSchedulerQueueInfo
- [YARN-4347](#) - Resource manager fails with Null pointer exception
- [YARN-4380](#) - TestResourceLocalizationService.testDownloadingResourcesOnContainerKill fails intermittently
- [YARN-4393](#) - Fix intermittent test failure for TestResourceLocalizationService#testFailedDirsResourceRelease
- [YARN-4573](#) - Fix test failure in TestRMAppTransitions#testAppRunningKill and testAppKilledKilled
- [YARN-4613](#) - Fix test failure in TestClientRMService#testGetClusterNodes
- [HBASE-14205](#) - RegionCoprocessorHost System.nanoTime() performance bottleneck
- [HBASE-14621](#) - ReplicationLogCleaner stuck on RS crash
- [HBASE-14923](#) - VerifyReplication should not mask the exception during result comparison
- [HBASE-14926](#) - Hung ThriftServer; no timeout on read from client; if client crashes, worker thread gets stuck reading
- [HBASE-15019](#) - Replication stuck when HDFS is restarted
- [HBASE-15031](#) - Fix merge of MVCC and SequenceID performance regression in branch-1.0
- [HBASE-15032](#) - hbase shell scan filter string assumes UTF-8 encoding
- [HBASE-15035](#) - bulkloading hfiles with tags that require splits do not preserve tags
- [HBASE-15052](#) - Use EnvironmentEdgeManager in ReplicationSource
- [HBASE-15104](#) - Occasional failures due to NotServingRegionException in IT tests
- [HBASE-15157](#) - Add *PerformanceTest for Append, CheckAnd*
- [HBASE-15213](#) - Fix increment performance regression caused by HBASE-8763 on branch-1.0

- [HIVE-7575](#) - GetTables thrift call is very slow
- [HIVE-10213](#) - MapReduce jobs using dynamic-partitioning fail on commit
- [HIVE-10514](#) - Fix MiniCliDriver tests failure
- [HIVE-11826](#) - 'hadoop.proxyuser.hive.groups' configuration does not prevent unauthorized user to access metastore
- [HIVE-11828](#) - beeline -f fails on scripts with tabs between column type and comment
- [HIVE-11977](#) - Hive should handle an external avro table with zero length files present
- [HIVE-12008](#) - Hive queries failing when using count(*) on column in view
- [HIVE-12388](#) - GetTables cannot get external tables when TABLE type argument is given
- [HIVE-12505](#) - Insert overwrite in same encrypted zone silently fails to remove some existing files
- [HIVE-12566](#) - Incorrect result returns when using COALESCE in WHERE condition with LEFT JOIN
- [HIVE-12713](#) - Miscellaneous improvements in driver compile and execute logging
- [HIVE-12784](#) - Group by SemanticException: Invalid column reference
- [HIVE-12790](#) - Metastore connection leaks in HiveServer2
- [HIVE-12795](#) - Vectorized execution causes ClassCastException
- [HIVE-12946](#) - alter table should also add default scheme and authority for the location similar to create table
- [HIVE-13039](#) - BETWEEN predicate is not functioning correctly with predicate pushdown on Parquet table
- [HIVE-13065](#) - Hive throws NPE when writing map type data to a HBase backed table
- [HUE-3106](#) - [filebrowser] Add support for full paths in zip file uploads
- [HUE-3110](#) - [oozie] Fix bundle submission when coordinator points to multiple bundles
- [HUE-3180](#) - [useradmin] Override duplicate username validation message
- [IMPALA-1702](#) - Check for duplicate table IDs at the end of analysis (issue not entirely fixed, but now fails gracefully)
- [IMPALA-2264](#) - Implicit casts to integers from decimals with higher precision sometimes allowed
- [IMPALA-2473](#) - Excessive memory usage by scan nodes
- [IMPALA-2621](#) - Fix flaky UNIX_TIMESTAMP() test
- [IMPALA-2643](#) - Nested inline view produces incorrect result when referencing the same column implicitly
- [IMPALA-2765](#) - AnalysisException: operands of type BOOLEAN and TIMESTAMP are not comparable when OUTER JOIN with CASE statement
- [IMPALA-2798](#) - After adding a column to avro table, Impala returns weird result if codegen is enabled.
- [IMPALA-2861](#) - Fix flaky scanner test added via IMPALA-2473 backport
- [IMPALA-2914](#) - Hit DCHECK Check failed: HasDateOrTime()
- [IMPALA-3034](#) - MemTracker leak on PHJ failure to spill
- [IMPALA-3085](#) - DataSinks' MemTrackers need to unregister themselves from parent
- [IMPALA-3093](#) - ReopenClient() could NULL out 'client_key' causing a crash
- [IMPALA-3095](#) - Allow additional Kerberos users to be authorized to access internal APIs
- [KITE-1114](#) - fix test
- [KITE-1114](#) - Fix missing license header
- [KITE-1114](#) - Kite CLI json-import HDFS temp file path not multiuser safe
- [OOZIE-2413](#) - Kerberos credentials can expire if the KDC is slow to respond
- [OOZIE-2428](#) - TestSLAService, TestSLAEventGeneration flaky tests
- [OOZIE-2432](#) - TestPurgeXCommand fails
- [OOZIE-2435](#) - TestCoordChangeXCommand is flaky
- [SENTRY-835](#) - Drop table leaves a connection open when using MetastoreListener
- [SENTRY-885](#) - DB name should be case insensitive in HDFS sync plugin
- [SENTRY-944](#) - Setting HDFS rules on Sentry managed hdfs paths should not affect original hdfs rules
- [SENTRY-953](#) - External Partitions which are referenced by more than one table can cause some unexpected behavior with Sentry HDFS sync
- [SENTRY-957](#) - Exceptions in MetastoreCacheInitializer should probably not prevent HMS from starting up
- [SENTRY-988](#) - It's better to let SentryAuthorization setter path always fall through and update HDFS
- [SENTRY-991](#) - backportRoles of Sentry Permission needs to be case insensitive
- [SENTRY-994](#) - SentryAuthorizationInfoX should override isSentryManaged

- [SENTRY-1002](#) - PathsUpdate.parsePath(path) will throw an NPE when parsing relative paths
- [SENTRY-1003](#) - backportSupport "reload" by updating the classpath of Sentry function aux jar path during run time
- [SENTRY-1008](#) - Path should be not be updated if the create/drop table/partition event fails
- [SENTRY-1044](#) - Tables with non-HDFS locations breaks HMS startup
- [SOLR-7281](#) - Add an overseer action to publish an entire node as 'down'
- [SOLR-8367](#) - Fix the LeaderInitiatedRecovery 'all replicas participate' fail-safe
- [SOLR-8371](#) - Try and prevent too many recovery requests from stacking up and clean up some faulty cancel recovery logic
- [SOLR-8372](#) - backportCanceled recovery can lead to data loss
- [SOLR-8575](#) - Addendum to Fix HDFSLogReader replay
- [SOLR-8575](#) - Fix HDFSLogReader replay status numbers and a performance bug where we can reopen FSDataInputStream too often
- [SOLR-8615](#) - Just like creating cores, we should use multiple threads when closing cores
- [SOLR-8720](#) - ZkController#publishAndWaitForDownStates should use #publishNodeAsDown
- [SOLR-8771](#) - Multithreaded core shutdown creates executor per core
- [SQOOP-2847](#) - Sqoop --incremental + missing parent --target-dir reports success with no data
- [SQOOP-2422](#) - Sqoop2: Test TestJSONIntermediateDataFormat is failing on JDK8
- [ZOOKEEPER-442](#) - Need a way to remove watches that are no longer of interest"

Issues Fixed in CDH 5.4.9

Known Issues Fixed

The following topics describe known issues fixed in CDH 5.4.9.

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera ("Cloudera Products"), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled "Arbitrary remote code execution with InvokerTransformer" and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Apache HBase

Data may not be replicated to worker cluster if multiwal multiplicity is set to greater than 1

Bug: [HBASE-13703](#), [HBASE-6617](#), [HBASE-14501](#).

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.4.9:

- [FLUME-2841](#) - Upgrade commons-collections to 3.2.2
- [HADOOP-7713](#) - dfs -count -q should label output column
- [HADOOP-11171](#) - Enable using a proxy server to connect to S3a
- [HADOOP-12568](#) - Update core-default.xml to describe posixGroups support
- [HADOOP-12577](#) - Bumped up commons-collections version to 3.2.2 to address a security flaw
- [HDFS-7785](#) - Improve diagnostics information for HttpPutFailedException
- [HDFS-7798](#) - Checkpointing failure caused by shared KerberosAuthenticator
- [HDFS-7871](#) - NameNodeEditLogRoller can keep printing 'Swallowing exception' message
- [HDFS-7990](#) - IBR delete ack should not be delayed
- [HDFS-8646](#) - Prune cached replicas from DatanodeDescriptor state on replica invalidation
- [HDFS-9123](#) - Copying from the root to a subdirectory should be forbidden
- [HDFS-9250](#) - Add Precondition check to LocatedBlock#addCachedLoc
- [HDFS-9273](#) - ACLs on root directory may be lost after NN restart
- [HDFS-9332](#) - Fix Precondition failures from NameNodeEditLogRoller while saving namespace
- [HDFS-9364](#) - Unnecessary DNS resolution attempts when creating NameNodeProxies
- [HDFS-9470](#) - Encryption zone on root not loaded from fsmage after NN restart
- [MAPREDUCE-6191](#) - Improve clearing stale state of Java serialization
- [MAPREDUCE-6549](#) - Multibyte delimiters with LineRecordReader cause duplicate records
- [YARN-4235](#) - FairScheduler PrimaryGroup does not handle empty groups returned for a user
- [HBASE-6617](#) - ReplicationSourceManager should be able to track multiple WAL paths
- [HBASE-12865](#) - WALs may be deleted before they are replicated to peers
- [HBASE-13134](#) - mutateRow and checkAndMutate APIs don't throw region level exceptions
- [HBASE-13618](#) - ReplicationSource is too eager to remove sinks
- [HBASE-13703](#) - ReplicateContext should not be a member of ReplicationSource
- [HBASE-14003](#) - Work around JDK-8044053
- [HBASE-14283](#) - Reverse scan doesn't work with HFile inline index/bloom blocks
- [HBASE-14374](#) - Backport parent 'HBASE-14317 Stuck FSHLog' issue to 1.1
- [HBASE-14501](#) - NPE in replication with TDE
- [HBASE-14533](#) - Connection Idle time 1 second is too short and the connection is closed too quickly by the ChoreService
- [HBASE-14547](#) - Add more debug/trace to zk-procedure
- [HBASE-14799](#) - Commons-collections object deserialization remote command execution vulnerability
- [HBASE-14809](#) - Grant / revoke Namespace admin permission to group
- [HIVE-10265](#) - Hive CLI crashes on != inequality
- [HIVE-11149](#) - Sometimes HashMap in PerfLogger.java hangs
- [HIVE-11616](#) - DelegationTokenSecretManager reuses the same objectstore, which has concurrency issues
- [HIVE-12058](#) - Change hive script to record errors when calling hbase fails
- [HIVE-12188](#) - DoAs does not work properly in non-Kerberos secured HS2
- [HIVE-12189](#) - The list in pushdownPreds of ppd.ExprWalkerInfo should not be allowed to grow very large
- [HIVE-12250](#) - ZooKeeper connection leaks in Hive's HBaseHandler
- [HIVE-12365](#) - Added resource path is sent to cluster as an empty string when externally removed
- [HIVE-12378](#) - Exception on HBaseSerDe.serialize binary field
- [HIVE-12406](#) - HIVE-9500 introduced incompatible change to LazySimpleSerDe public interface
- [HIVE-12418](#) - HiveHBaseTableInputFormat.getRecordReader() causes ZooKeeper connection leak
- [HUE-2941](#) - [hadoop] Cache the active RM HA
- [HUE-3035](#) - [beeswax] Optimize sample data query for partitioned tables
- [IMPALA-1459](#) - Fix migration/assignment of On-clause predicates inside inline views

- [IMPALA-1675](#) - Avoid overflow when adding large intervals to TIMESTAMPs
- [IMPALA-1746](#) - QueryExecState does not check for query cancellation or errors
- [IMPALA-1917](#) - Do not register aux equivalence predicates with NULL on either side
- [IMPALA-1949](#) - Analysis exception when a binary operator contain an IN operator with
- [IMPALA-2086/IMPALA-2090](#) - Avoid boost year/month interval logic
- [IMPALA-2141](#) - UnionNode::getNext() does not check for query errors
- [IMPALA-2252](#) - Crash (likely race) tearing down BufferedBlockMgr on query failure
- [IMPALA-2260](#) - Adding a large hour interval caused an interval overflow
- [IMPALA-2265](#) - Sorter was not checking the returned Status of PrepareRead
- [IMPALA-2273](#) - Make MAX_PAGE_HEADER_SIZE configurable
- [IMPALA-2286](#) - Fix race between ~BufferedBlockMgr() and BufferedBlockMgr::Create()
- [IMPALA-2344](#) - Work-around IMPALA-2344 Fail query with OOM in case block->Pin() fails
- [IMPALA-2357](#) - Fix spilling sorts with var-len slots that are NULL or empty
- [IMPALA-2446](#) - Fix wrong predicate assignment in outer joins
- [IMPALA-2533](#) - Impala throws IllegalStateException when inserting data into a partition
- [IMPALA-2559](#) - Fix check failed: sorter_runs_.back()->is_pinned_
- [IMPALA-2664](#) - Avoid sending large partition stats objects over thrift
- [IMPALA-2731](#) - Refactor MemPool usage in HBase scan node
- [KITE-1089](#) - readAvroContainer morphline command should work even if the Avro writer schema of each input file is different
- [PIG-3641](#) - Split "otherwise" producing incorrect output when combined with ColumnPruning
- [SENTRY-565](#) - Improve performance of filtering Hive SHOW commands
- [SENTRY-702](#) - Hive binding should support RELOAD command
- [SENTRY-936](#) - getGroup and getUser should always return original hdfs values for paths in prefixes which are not Sentry managed
- [SENTRY-960](#) - Blacklist reflect, java_method using hive.server2.builtin.udf.blacklist
- [SOLR-6443](#) - backportDisable test that fails on Jenkins with SolrCore.getOpenCount()==2
- [SOLR-7049](#) - LIST Collections API call should be processed directly by the CollectionsHandler instead of the OverseerCollectionProcessor
- [SOLR-7552](#) - Support using ZkCredentialsProvider/ZkACLProvider in custom filter
- [SOLR-7989](#) - After a new leader is elected, it should ensure it's state is ACTIVE if it has already registered with ZK
- [SOLR-8075](#) - Leader Initiated Recovery should not stop a leader that participated in an election with all of it's replicas from becoming a valid leader
- [SOLR-8223](#) - Avoid accidentally swallowing OutOfMemoryError
- [SOLR-8288](#) - DistributedUpdateProcessor#doFinish should explicitly check and ensure it does not try to put itself into LIR
- [SPARK-11484](#) - [WEBUI] Using proxyBase set by spark AM
- [SPARK-11652](#) - [CORE] Remote code execution with InvokerTransformer

Issues Fixed in CDH 5.4.8

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.4.8:

- [FLUME-2095](#) - JMS source with TIBCO
- [HADOOP-11261](#) - Set custom endpoint for S3A
- [HADOOP-12404](#) - Disable caching for JarURLConnection to avoid sharing JarFile with other users when loading resource from URL in Configuration class
- [HADOOP-12413](#) - AccessControlList should avoid calling getGroupNames in isUserInList with empty groups
- [HDFS-7916](#) - 'reportBadBlocks' from datanodes to standby Node BPSERVICEACTOR goes for infinite loop
- [HDFS-7978](#) - Add LOG.isDebugEnabled() guard for some LOG.debug()
- [HDFS-8384](#) - Allow NN to startup if there are files having a lease but are not under construction

- [HDFS-8735](#) - Inotify: All events classes should implement toString() API
- [HDFS-8860](#) - Remove unused Replica copyOnWrite code
- [HDFS-8964](#) - When validating the edit log, do not read at or beyond the file offset that is being written
- [HDFS-8965](#) - Harden edit log reading code against out of memory errors
- [MAPREDUCE-5918](#) - LineRecordReader can return the same decompressor to CodecPool multiple times
- [MAPREDUCE-5948](#) - org.apache.hadoop.mapred.LineRecordReader does not handle multibyte record delimiters well
- [MAPREDUCE-6273](#) - HistoryFileManager should check whether summaryFile exists to avoid FileNotFoundException causing HistoryFileInfo into MOVE_FAILED state
- [MAPREDUCE-6481](#) - LineRecordReader may give incomplete record and wrong position/key information for uncompressed input sometimes
- [MAPREDUCE-6484](#) - YARN Client uses local address instead of RM address as token renewer in a secure cluster when RM HA is enabled
- [YARN-2666](#) - TestFairScheduler.testContinuousScheduling fails intermittently
- [YARN-3385](#) - Fixed a race-condition in ResourceManager's ZooKeeper based state-store to avoid crashing on duplicate deletes
- [YARN-3469](#) - ZKRMStateStore: Avoid setting watches that are not required.
- [YARN-3943](#) - Use separate threshold configurations for disk-full detection and disk-not-full detection
- [HBASE-13217](#) - Procedure fails due to ZK issue
- [HBASE-13331](#) - Exceptions from DFS client can cause CatalogJanitor to delete referenced files
- [HBASE-13388](#) - Handling NullPointerException in ZKProcedureMemberRpcs while getting ZNode data
- [HBASE-13933](#) - DBE's seekBefore with tags corrupts the tag's offset information thus leading to incorrect results
- [HBASE-14196](#) - Thrift server idle connection timeout issue
- [HBASE-14302](#) - TableSnapshotInputFormat should not create back references when restoring snapshot
- [HBASE-14347](#) - Add a switch to DynamicClassLoader to disable it
- [HBASE-14385](#) - Close the sockets that is missing in connection closure
- [HBASE-14394](#) - Properly close the connection after reading records from table
- [HBASE-14471](#) - Thrift - HTTP Error 413 full HEAD if using Kerberos authentication
- [HBASE-14492](#) - Increase REST server header buffer size from 8k to 64k
- [HIVE-5545](#) - HCatRecord getInteger method returns String when used on Partition columns of type INT
- [HIVE-8529](#) - HiveSessionImpl#fetchResults should not try to fetch operation log when hive.server2.logging.operation.enabled is false
- [HIVE-9867](#) - Migrate usage of deprecated Calcite methods
- [HIVE-9984](#) - JoinReorder's getOutputSize is exponential
- [HIVE-10021](#) - "Alter index rebuild" statements submitted through HiveServer2 fail when Sentry is enabled
- [HIVE-10122](#) - Hive metastore filter-by-expression is broken for non-partition expressions
- [HIVE-10421](#) - DROP TABLE with qualified table name ignores database name when checking partitions
- [HIVE-10451](#) - PTF deserializer fails if values are not used in reducer
- [HIVE-10658](#) - Insert with values clause may expose data that should be encrypted
- [HIVE-10980](#) - Merge of dynamic partitions loads all data to default partition
- [HIVE-11077](#) - Part of Exchange partition does not properly populate fields for post/pre execute hooks.
- [HIVE-11440](#) - Create Parquet predicate push down (PPD) unit tests and q-tests
- [HIVE-11504](#) - Predicate pushing down does not work for float type for Parquet
- [HIVE-11590](#) - AvroDeserializer is very chatty
- [HIVE-11618](#) - BackportCorrect the SARG api to reunify the PredicateLeaf.Type INTEGER and LONG
- [HIVE-11657](#) - HIVE-2573 introduces some issues during metastore init (and CLI init)
- [HIVE-11695](#) - If user has no permission to create LOCAL DIRECTORY, the Hql does not throw any exception and fails silently
- [HIVE-11696](#) - Exception when table-level serde is Parquet while partition-level serde is JSON
- [HIVE-11712](#) - Duplicate groupby keys cause ClassCastException
- [HIVE-11737](#) - IndexOutOfBounds compiling query with duplicated groupby keys

- [HIVE-11745](#) - Alter table Exchange partition with multiple partition_spec is not working
- [HIVE-11816](#) - Upgrade groovy to 2.4.4
- [HIVE-11824](#) - Insert to local directory causes staging directory to be copied
- [HIVE-11843](#) - Add 'sort by c' to Parquet PPD q-tests to avoid different output issues with hadoop-1
- [HIVE-11891](#) - Add basic performance logging to metastore calls
- [HIVE-11926](#) - Backport:Stats annotation might not extract stats for varchar/decimal columns
- [HIVE-11982](#) - Some test cases for union all fail with recent changes
- [HIVE-11995](#) - Remove repetitively setting permissions in insert/load overwrite partition
- [HUE-2881](#) - [oozie] A fork can point to a deleted node
- [IMPALA-1136](#) - Support loading Avro tables without an explicit Avro schema
- [IMPALA-1899](#) - Cleanup handling of Hive's field schema3e0fee5 IMPALA-2369, IMPALA-2435: Impala crashes when the sorter hits an OOM error
- [IMPALA-2130](#) - Wrong verification of Parquet file version
- [IMPALA-2161](#) - Skip \u0000 characters when dealing Avro schemas
- [IMPALA-2165](#) - Avoid cardinality 0 in scan nodes of small tables and low selectivity
- [IMPALA-2168](#) - Do not try to access streams of repartitioned spilled partition in right-joins
- [IMPALA-2213](#) - Make Parquet scanner fail query if the file size metadata is stale
- [IMPALA-2249](#) - Avoid allocating StringBuffer > 1GB in ScannerContext::Stream::GetBytesInternal()
- [IMPALA-2256](#) - Handle joins with right side of high cardinality and zero materialized slots
- [IMPALA-2270](#) - Avoid FnvHash64to32 with empty inputs
- [IMPALA-2284](#) - Disallow long (1<<30) strings in group_concat()
- [IMPALA-2292](#) - Change the type of timestamp_col to string in the table no_avro_schema.
- [IMPALA-2314](#) - LargestSpilledPartition was not checking if partition is closed
- [IMPALA-2348](#) - The catalog does not close the connection to HMS during table invalidation
- [IMPALA-2364](#) - Wrong DCHECK in PHJ::ProcessProbeBatch
- [IMPALA-2366](#) - Check fread return code correctly
- [IMPALA-2440](#) - Fix old HJ full outer join with no rows
- [IMPALA-2477](#) - Parquet metadata randomly 'appears stale'
- [IMPALA-2514](#) - DCHECK on destroying an ExprContext
- [KITE-1069](#) - Make zkClientSessionTimeout and zkClientConnectTimeout configurable in SolrLocator
- [KITE-1074](#) - Partial updates aka Atomic updates with loadSolr aren't recognized with SolrCloud
- [MAHOUT-1771](#) - Cluster dumper omits indices and 0 elements for dense vector or sparse containing 0s
- [OOZIE-2376](#) - Default action configs not honored if no <configuration> section in workflow
- [SENTRY-878](#) - collect_list missing from HIVE_UDF_WHITE_LIST
- [SENTRY-884](#) - Give execute permission by default to paths managed by sentry
- [SENTRY-893](#) - Synchronize calls in SentryClient and create sentry client once per request in SimpleDBProvider
- [SOLR-5776](#) - Use less TLS/SSL in a test run
- [SOLR-7109](#) - Indexing threads stuck during network partition can put leader into down state
- [SOLR-7844](#) - Zookeeper session expiry during shard leader election can cause multiple leaders
- [SOLR-7956](#) - There are interrupts on shutdown in places that can cause ChannelAlreadyClosed exceptions which prevents proper closing of transaction logs and can poison the IndexWriter and interfere with the HDFS client
- [SOLR-8046](#) - HdfsCollectionsAPIDistributedZkTest checks that no transaction logs failed to be opened during the test but does not isolate this to the test and could fail due to other tests
- [SOLR-8069](#) - Ensure that only the valid ZooKeeper registered leader can put a replica into Leader Initiated Recovery
- [SOLR-8075](#) - Leader Initiated Recovery should not stop a leader that participated in an election with all of it's replicas from becoming a valid leader
- [SOLR-8077](#) - Replication can still cause index corruption
- [SOLR-8085](#) - Fix a variety of issues that can result in replicas getting out of sync
- [SOLR-8094](#) - HdfsUpdateLog should not replay buffered documents as a replacement to dropping them
- [SOLR-8095](#) - Add enable prop for HDFS Locality Metrics

- [SOLR-8121](#) - It looks like ChaosMonkeySafeLeader test can fail with replica inconsistency because `waitForThingsToLevelOut` can pass while state is still changing.
- [SPARK-6880](#) - Spark Shutdowns with `NoSuchElementException` when running parallel collect on cachedRDD
- [SQOOP-2597](#) - Missing method `AvroSchemaGenerator.generate()`

Published Known Issues Fixed

As a result of the above fixes, the following issues, previously published as [Known Issues in CDH 5](#) on page 1944, are also fixed.

Spurious warning in MRv1 jobs

The `mapreduce.client.genericoptionsparser.used` property is not correctly checked by `JobClient` and this leads to a spurious warning.

Cloudera Bug: CDH-9740

Workaround: MapReduce jobs using `GenericOptionsParser` or implementing `Tool` can remove the warning by setting this property to true.

Spark Sink requires spark-assembly.jar in Flume classpath

In CDH 5.4.0, Flume requires `spark-assembly.jar` in the Flume classpath to use the Spark Sink. Without this, the sink fails with a dependency issue.

Bug: [SPARK-7038](#)

Cloudera Bug: CDH-27210

Workaround: Use the Spark Sink from CDH 5.3 with Spark from CDH 5.4, or add `spark-assembly.jar` to the `FLUME_CLASSPATH`.

Streaming incompatibility between Spark 1.2 and 1.3

Applications built as a JAR with dependencies ("uber JAR") must be built for the specific version of Spark running on the cluster.

Cloudera Bug: CDH-26527

Workaround: Rebuild the JAR with the Spark dependencies in `pom.xml` pointing to the specific version of Spark running on the target cluster.

Configuring more than one NT domain does not work in CDH 5.4.0

Trying to add users and groups using the multi-NT domain feature (<http://gethue.com/hadoop-tutorial-make-hadoop-more-accessible-by-integrating-multiple-ldap-servers/>) produces an error.

Bug: [HUE-2665](#)

Cloudera Bug: CDH-26431

Workaround: None.

If Sentry is enabled, the RELOAD command cannot be executed in the Hive CLI or Beeline.

Bug: [SENTRY-702](#)

Cloudera Bug: CDH-25786

Workaround: None.

Issues Fixed in CDH 5.4.7

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.4.7:

- [CRUNCH-525](#) - Correct (more) accurate default scale factors for built-in MapFn implementations
- [CRUNCH-527](#) - Use hash smearing for partitioning

- [CRUNCH-528](#) - Improve Pair comparison
- [CRUNCH-531](#) - Fix split graph rendering typo.
- [CRUNCH-535](#) - call initCredentials on the job
- [CRUNCH-536](#) - Refactor CrunchControlledJob.Hook interface and make it client-accessible
- [CRUNCH-539](#) - Fix reading WritableComparables bimap
- [CRUNCH-540](#) - Make AvroReflectDeepCopier serializable
- [CRUNCH-543](#) - Have AvroPathPerKeyTarget handle child directories properly
- [CRUNCH-544](#) - Improve performance/serializability of materialized toMap.
- [CRUNCH-546](#) - Remove calls to CellUtil.cloneXXX
- [CRUNCH-547](#) - Properly handle nullability for Avro union types
- [CRUNCH-548](#) - Have the AvroReflectDeepCopier use the class of the source object when constructing new instances instead of the target class
- [CRUNCH-551](#) - Make the use of Configuration objects consistent in CrunchInputSplit and CrunchRecordReader
- [CRUNCH-553](#) - Fix record drop issue that can occur w/From.formattedFile TableSources
- [FLUME-1934](#) - Spooling Directory Source dies on encountering zero-byte files.
- [FLUME-2753](#) - Error when specifying empty replace string in Search and Replace Interceptor
- [HADOOP-12317](#) - Applications fail on NM restart on some linux distro because NM container recovery declares AM container as LOST
- [HDFS-8806](#) - Inconsistent metrics: number of missing blocks with replication factor 1 not properly cleared
- [HDFS-8850](#) - VolumeScanner thread exits with exception if there is no block pool to be scanned but there are suspicious blocks.
- [MAPREDUCE-5817](#) - Mappers get rescheduled on node transition even after all reducers are completed.
- [MAPREDUCE-6277](#) - Job can post multiple history files if attempt loses connection to the RM
- [MAPREDUCE-6439](#) - AM may fail instead of retrying if RM shuts down during the allocate call.
- [YARN-2921](#) - Fix MockRM/MockAM#waitForState sleep too long.
- [YARN-3823](#) - Fix mismatch in default values for yarn.scheduler.maximum-allocation-vcores property
- [YARN-3990](#) - AsyncDispatcher may overloaded with RMAppNodeUpdateEvent when Node is connected/disconnected
- [HBASE-13329](#) - ArrayIndexOutOfBoundsException in CellComparator#getMinimumMidpointArray.
- [HBASE-13437](#) - ThriftServer leaks ZooKeeper connections
- [HBASE-13471](#) - Fix a possible infinite loop in doMiniBatchMutation
- [HBASE-13684](#) - Allow mlockagent to be used when not starting as root
- [HBASE-14162](#) - Fixing maven target for regenerating thrift classes fails against 0.9.2
- [HBASE-14354](#) - Minor improvements for usage of the mlock agent
- [HIVE-7476](#) - CTAS does not work properly for s3
- [HIVE-9327](#) - CBO (Calcite Return Path): Removing Row Resolvers from ParseContext
- [HIVE-9512](#) - HIVE-9327 causing regression in stats annotation
- [HIVE-9580](#) - Server returns incorrect result from JOIN ON VARCHAR columns
- [HIVE-9613](#) - Left join query plan outputs wrong column when using subquery
- [HIVE-10085](#) - Lateral view on top of a view throws RuntimeException
- [HIVE-10140](#) - Window boundary is not compared correctly
- [HIVE-10288](#) - Cannot call permanent UDFs
- [HIVE-10319](#) - Hive CLI startup takes a long time with a large number of databases
- [HIVE-10719](#) - Hive metastore failure when alter table rename is attempted.
- [HIVE-10875](#) - Select query with view in subquery adds underlying table as direct input
- [HIVE-10906](#) - Value based UDAF function without orderby expression throws NPE
- [HIVE-10911](#) - Add support for date datatype in the value based windowing function
- [HIVE-10972](#) - DummyTxnManager always locks the current database in shared mode, which is incorrect
- [HIVE-10985](#) - Value based windowing on timestamp and double can't handle NULL value
- [HIVE-10996](#) - Aggregation / Projection over Multi-Join Inner Query producing incorrect results
- [HIVE-11139](#) - PROPOSEDQTest combine2_hadoop20.q fails when using -Phadoop-1 profile due to

- [HIVE-11172](#) - Vectorization wrong results for aggregate query with where clause without group by
- [HIVE-11203](#) - Beeline force option does not force execution when errors occurred in a script.
- [HIVE-11250](#) - Change in spark.executor.instances (and others) does not take effect after RSC is launched for HS2
- [HIVE-11255](#) - `get_table_objects_by_name()` in `HiveMetaStore.java` needs to retrieve table objects in multiple batches
- [HIVE-11258](#) - The function `drop_database_core()` of `HiveMetaStore.java` may not drop all the tables
- [HIVE-11271](#) - `java.lang.IndexOutOfBoundsException` when union all with if function
- [HIVE-11288](#) - Avro SerDe InstanceCache returns incorrect schema
- [HIVE-11333](#) - ColumnPruner prunes columns of UnionOperator that should be kept
- [HIVE-11502](#) - Map side aggregation is extremely slow
- [HIVE-11604](#) - HIVE return wrong results in some queries with PTF function
- [HIVE-11620](#) - Fix several qtest output order
- [HUE-2873](#) - [oozie] Handle TransactionManagementError on workflow dashboard
- [HUE-2877](#) - [desktop] Add pyasn1 and ndg_httpsclient to support SSL Server Name Indication
- [HUE-2880](#) - [hadoop] Fix uploading large files to a kerberized HTTPFS
- [HUE-2882](#) - [oozie] Fix parsing error when workflow job uses Australian timezone
- [HUE-2883](#) - [impala] Canceling a query shows an error message
- [HUE-2885](#) - [oozie] Java options java-opts not generated correctly in XML
- [HUE-2893](#) - [desktop] Backport CherryPy SSL file upload fix
- [HUE-2903](#) - [oozie] Fix error with Workflow parameter on rerun
- [IMPALA-1737](#) - Substitute an InsertStmt's partition key exprs with the root node's smap.
- [IMPALA-1756](#) - Constant filter expressions are not checked for errors and state cleanup is not done before throwing exception.
- [IMPALA-1898](#) - Explicit aliases + ordinals analysis bug
- [IMPALA-1983](#) - Warn if table stats are potentially corrupt.
- [IMPALA-1987](#) - Fix TupleIsNullPredicate to return false if no tuples are nullable.
- [IMPALA-2088](#) - Fix planning of empty union operands with analytics.
- [IMPALA-2089](#) - Retain eq predicates bound by grouping slots with complex grouping exprs.
- [IMPALA-2178](#) - fix `Expr::ComputeResultsLayout()` logic.
- [IMPALA-2199](#) - Row count not set for empty partition when spec is used with compute incremental stats
- [IMPALA-2201](#) - Unconditionally update the partition stats and row count.
- [IMPALA-2203](#) - Set an InsertStmt's result exprs from the source statement's result exprs.
- [IMPALA-2216](#) - Set the output smap of an EmptySetNode produced from an empty inline view.
- [IMPALA-2239](#) - update misc.test to match the new .test file format.
- [IMPALA-2266](#) - Pass correct child node in 2nd phase merge aggregation.
- [KITE-1053](#) - Fix int overflow bug in FS writer.
- [SENTRY-810](#) - CTAS without location is not verified properly
- [SOLR-7135](#) - Allow the server build.xml 'sync-hack' target to be skipped by specifying a system property.
- [SOLR-7999](#) - SolrRequetParserTest#testStreamURL started failing.
- [SPARK-8606](#) - Prevent exceptions in `RDD.getPreferredLocations()` from crashing DAGScheduler
- [ZOOKEEPER-442](#) - need a way to remove watches that are no longer of interest

Issues Fixed in CDH 5.4.5

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.4.5:

- [CRUNCH-508](#) - Improve performance of Scala Enumeration counters in Scrunch
- [CRUNCH-511](#) - Scrunch product type support should use `derived()` instead of `derivedImmutable()`
- [CRUNCH-514](#) - AvroDerivedDeepCopier should initialize delegate MapFns
- [CRUNCH-516](#) - Scrunch needs some additional null checks
- [CRUNCH-530](#) - Fix object reuse bug in `GenericRecordToTuple`

- [CRUNCH-542](#) - Wider tolerance for flaky scrunch PCollectionTest
- [FLUME-2215](#) - ResettableFileInputStream can't support ucs-4 character
- [FLUME-2732](#) - Make maximum tolerated failures before shutting down and recreating client in AsyncHbaseSink configurable
- [FLUME-2738](#) - Async HBase sink FD leak on client shutdown
- [FLUME-2749](#) - Kerberos configuration error when using short names in multiple HDFS Sinks
- [HADOOP-12017](#) - Hadoop archives command should use configurable replication factor when closing
- [HADOOP-12103](#) - Small refactoring of DelegationTokenAuthenticationFilter to allow code sharing
- [HADOOP-8151](#) - Error handling in snappy decompressor throws invalid exceptions
- [HDFS-7501](#) - TransactionsSinceLastCheckpoint can be negative on SBNS
- [HDFS-7546](#) - Document, and set an accepting default for dfs.namenode.kerberos.principal.pattern
- [HDFS-7890](#) - Improve information on Top users for metrics in RollingWindowsManager and lower log level
- [HDFS-7894](#) - Rolling upgrade readiness is not updated in jmx until query command is issued.
- [HDFS-8072](#) - Reserved RBW space is not released if client terminates while writing block
- [HDFS-8337](#) - Accessing httpfs via webhdfs doesn't work from a jar with kerberos
- [HDFS-8656](#) - Preserve compatibility of ClientProtocol#rollingUpgrade after finalization
- [HDFS-8681](#) - BlockScanner is incorrectly disabled by default
- [MAPREDUCE-5965](#) - Hadoop streaming throws error if list of input files is high.
- [YARN-3143](#) - RM Apps REST API can return NPE or entries missing id and other fields
- [YARN-3453](#) - Fair Scheduler: Parts of preemption logic uses DefaultResourceCalculator even in DRF mode causing thrashing
- [YARN-3535](#) - Scheduler must re-request container resources when RMContainer transitions from ALLOCATED to KILLED
- [YARN-3793](#) - Several NPEs when deleting local files on NM recovery
- [YARN-3842](#) - NMProxy should retry on NMNotYetReadyException
- [HBASE-13342](#) - Fix incorrect interface annotations
- [HBASE-13419](#) - Thrift gateway should propagate text from exception causes.
- [HBASE-13491](#) - Fix bug in FuzzyRowFilter#getNextForFuzzyRule
- [HBASE-13851](#) - RpcClientImpl.close() can hang with cancelled replica RPCs
- [HBASE-13885](#) - ZK watches leaks during snapshots
- [HBASE-13958](#) - RESTApiClusterManager calls kill() instead of suspend() and resume()
- [HBASE-13995](#) - ServerName is not fully case insensitive
- [HBASE-14027](#) - Clean up netty dependencies
- [HBASE-14045](#) - Bumping thrift version to 0.9.2.
- [HBASE-14076](#) - ResultSerialization and MutationSerialization can throw InvalidProtocolBufferException when serializing a cell larger than 64MB
- [HIVE-10252](#) - Make PPD work for Parquet in row group level
- [HIVE-10270](#) - Cannot use Decimal constants less than 0.1BD
- [HIVE-10553](#) - Remove hardcoded Parquet references from SearchArgumentImpl SearchArgumentImpl
- [HIVE-10706](#) - Make vectorized_timestamp_funcs test more stable
- [HIVE-10801](#) - 'drop view' fails throwing java.lang.NullPointerException
- [HIVE-10808](#) - Inner join on Null throwing Cast Exception
- [HIVE-11150](#) - Remove wrong warning message related to chgrp
- [HIVE-11174](#) - Hive does not treat floating point signed zeros as equal
- [HIVE-11216](#) - UDF GenericUDFMapKeys throws NPE when a null map value is passed in
- [HIVE-11401](#) - Predicate push down does not work with Parquet when partitions are in the expression expression
- [HIVE-6099](#) - Multi insert does not work properly with distinct count
- [HIVE-9500](#) - Support nested structs over 24 levels
- [HIVE-9665](#) - Parallel move task optimization causes race condition
- [HIVE-10427](#) - collect_list() and collect_set() should accept struct types as argument

- [HIVE-10437](#) - NullPointerException on queries where map/reduce is not involved on tables with partitions
- [HIVE-10895](#) - ObjectStore does not close Query objects in some calls, causing a potential leak in some metastore db resources
- [HIVE-10976](#) - Redundant HiveMetaStore connect check in HS2 CLIService start
- [HIVE-10977](#) - No need to instantiate MetaStoreDirectSql when HMS DirectSql is disabled
- [HIVE-11095](#) - Fix SerDeUtils bug when Text is reused
- [HIVE-11100](#) - Beeline should escape semi-colon in queries
- [HIVE-11112](#) - ISO-8859-1 text output has fragments of previous longer rows appended
- [HIVE-11157](#) - Hive.get(HiveConf) returns same Hive object to different user sessions
- [HIVE-11194](#) - Exchange partition on external tables should fail with error message when target folder already exists
- [HIVE-11433](#) - NPE for a multiple inner join query
- [HIVE-9767](#) - Fixes in Hive UDF to be usable in Pig
- [HIVE-10629](#) - Dropping table in an encrypted zone does not drop warehouse directory
- [HIVE-10630](#) - Renaming tables across encryption zones renames table even though the operation throws error
- [HIVE-10659](#) - Beeline command which contains semi-colon as a non-command terminator will fail
- [HIVE-10788](#) - Change sort_array to support non-primitive types
- [HIVE-10895](#) - ObjectStore does not close Query objects in some calls causing potential leak.
- [HIVE-11109](#) - Replication factor is not properly set in SparkHashTableSinkOperator [Spark Branch]
- [HIVE-10594](#) - Remote Spark client doesn't use Kerberos keytab to authenticate [Spark Branch]
- [HUE-2618](#) - [hive] Recent query results show character encoding in view
- [HUE-2767](#) - [impala] Issue showing sample data for a table
- [HUE-2796](#) - sync_groups_on_login doesn't work with posixGroups
- [HUE-2807](#) - [useradmin] Support deleting numeric groups
- [HUE-2808](#) - [dbquery] Add row numbers to support default order by
- [HUE-2813](#) - [hive] Report when Hue server is down when trying to execute a query
- [HUE-2814](#) - Revert pyopenssl 0.13.1
- [HUE-2835](#) - Fixed issue with DN's that have weird comma location
- [HUE-2840](#) - [useradmin] Fix create home directories for Add/Sync LDAP group
- [HUE-2849](#) - [useradmin] Fix exception in Add/Sync LDAP group for undefined group name
- [IMPALA-1929](#) - Avoiding a DCHECK of NULL hash table in spilled right joins
- [IMPALA-2136](#) - Bug in PrintTColumnValue caused wrong stats for TINYINT partition cols
- [IMPALA-2133](#) - Properly unescape string value for HBase filters
- [IMPALA-2018](#) - Where clause does not propagate to joins inside nested views
- [IMPALA-2064](#) - Add effective_user() builtin
- [IMPALA-2125](#) - Make UTC to local TimestampValue conversion faster.
- [IMPALA-2048](#) - Impala DML/DDDL operations corrupt table metadata leading to Hive query failures
- [KITE-1014](#) - Fix support for Hive datasets on Kerberos enabled clusters.
- [KITE-1015](#) - Add "replaceValues" morphline command that replaces all matching record field values with a given replacement string
- [KITE-462](#) - Oozie jobs do not pass credentials
- [KITE-976](#) - DatasetKeyInputFormat/DatasetKeyOutputFormat not setting job configuration before loading dataset
- [KITE-1030](#) - readCSV WARN log msg on overly long lines where quoteChar is non-empty should print the whole record seen so far
- [OOZIE-2268](#) - Update ActiveMQ version for security and other fixes
- [OOZIE-2286](#) - Update Log4j and Log4j-extras to latest 1.2.x release
- [PIG-4053](#) - PIG-4053: TestMRCompiler succeeded with sun jdk 1.6 while failed with sun jdk 1.7
- [PIG-4338](#) - PIG-4338: Fix test failures with JDK8
- [PIG-4326](#) - PIG-4326: AvroStorageSchemaConversionUtilities does not properly convert schema for maps of arrays of records

- [SENTRY-695](#) - Sentry service should read the hadoop group mapping properties from core-site
- [SENTRY-721](#) - HDFS Cascading permissions not applied to child file ACLs if a direct grant exists
- [SENTRY-752](#) - Sentry service audit log file name format should be consistent
- [SQLR-7457](#) - Make DirectoryFactory publishing MBeanInfo extensible
- [SQLR-7458](#) - Expose HDFS Block Locality Metrics
- [SPARK-6480](#) - histogram() bucket function is wrong in some simple edge cases
- [SPARK-6954](#) - ExecutorAllocationManager can end up requesting a negative number of executors
- [SPARK-7503](#) - Resources in .sparkStaging directory can't be cleaned up on error
- [SPARK-7705](#) - Cleanup of .sparkStaging directory fails if application is killed
- [SQOOP-2103](#) - Not able define Decimal(n,p) data type in map-column-hive option
- [SQOOP-2149](#) - Update Kite dependency to 1.0.0
- [SQOOP-2252](#) - Add default to Avro Schema
- [SQOOP-2294](#) - Change to Avro schema name breaks some use cases
- [SQOOP-2295](#) - Hive import with Parquet should append automatically
- [SQOOP-2327](#) - Sqoop2: Change package name from Authorization to authorization
- [SQOOP-2339](#) - Move sub-directory might fail in append mode
- [SQOOP-2362](#) - Add oracle direct mode in list of supported databases
- [SQOOP-2400](#) - hive.metastore.sasl.enabled should be set to true for Oozie integration
- [SQOOP-2406](#) - Add support for secure mode when importing Parquet files into Hive
- [SQOOP-2437](#) - Use hive configuration to connect to secure metastore

Issues Fixed in CDH 5.4.4

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.4.4:

- [HIVE-10572](#) - Improve Hive service test to check empty string
- [HIVE-9934](#) - Vulnerability in LdapAuthenticationProviderImpl enables HiveServer2 client to degrade the authentication mechanism to "none", allowing authentication without password
- [HIVE-10006](#) - RSC has memory leak while execute multi queries.
- [HUE-2814](#) - Revert pyopenssl 0.13.1

Published Known Issues Fixed

As a result of the above fixes, the following issues, previously published as [Known Issues in CDH 5](#) on page 1944, are also fixed.

Hue with TLS/SSL Fails to Start in CDH 5.4.3

In CDH 5.4.3, Hue with TLS/SSL fails to start because a pyOpenSSL package is missing in the parcel. This applies to both new installs and upgrades and is not operating-system specific.

Bug: CDH-29076

Release affected: CDH 5.4.3

Release containing the fix: CDH 5.4.4

Workaround:

1. Download the package:

a. `cd /tmp`

b. `curl -O https://pypi.python.org/packages/source/p/pyOpenSSL/pyOpenSSL-0.13.tar.gz`

2. Determine the Hue installation directory:

- Parcels:

```
export HUE_DIR=/opt/cloudera/parcels/CDH-5.4.3-1.cdh5.4.3.p0.6/lib/hue
```

- Packages:

```
export HUE_DIR=/usr/lib/hue
```

3. Change to the Hue installation directory:

```
cd $HUE_DIR
```

4. Do the following, depending on your OS:

- On CentOS/RedHat 6.x:
 1. `sudo yum install gcc python-devel openssl-devel`
 2. `sudo ./build/env/bin/python ./build/env/bin/pip -v install /tmp/pyOpenSSL-0.13.tar.gz`
- On Ubuntu 14.04:
 1. `sudo apt-get install gcc python-dev python-pip libssl-dev`
 2. `sudo pip install --target=`pwd`/`ls -d build/env/lib/python*/site-packages` /tmp/pyOpenSSL-0.13.tar.gz`
- On other platforms, [contact Support](#) for assistance.

Issues Fixed in CDH 5.4.3

Upgrades to CDH 5.4.1 from Releases Earlier than 5.4.0 May Fail

Problem: Because of a change in the implementation of the NameNode metadata upgrade mechanism, upgrading to CDH 5.4.1 from a version lower than 5.4.0 can take an inordinately long time. In a cluster with NameNode high availability (HA) configured and a large number of edit logs, the upgrade can fail, with errors indicating a timeout in the pre-upgrade step on JournalNodes.

What to do:

To avoid the problem: Do not upgrade to CDH 5.4.1; upgrade to CDH 5.4.2 instead.

If you experience the problem: If you have already started an upgrade and seen it fail, contact Cloudera Support. This problem involves no risk of data loss, and manual recovery is possible.

If you have already completed an upgrade to CDH 5.4.1, or are installing a new cluster: In this case you are not affected and can continue to run CDH 5.4.1.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.4.3:

- [HADOOP-12043](#) - Display warning if defaultFs is not set when running fs commands.
- [HADOOP-11969](#) - ThreadLocal initialization in several classes is not thread safe
- [HADOOP-11402](#) - Negative user-to-group cache entries are never cleared for never-again-accessed users
- [HADOOP-11238](#) - Update the NameNode's Group Cache in the background when possible
- [HDFS-8535](#) - Clarify that dfs usage in dfsadmin -report output includes all block replicas.
- [HDFS-8486](#) - DN startup may cause severe data loss
- [HDFS-7917](#) - Use file to replace data dirs in test to simulate a disk failure.
- [HDFS-7833](#) - DataNode reconfiguration does not recalculate valid volumes required, based on configured failed volumes tolerated.
- [HDFS-7604](#) - Track and display failed DataNode storage locations in NameNode.
- [HDFS-8380](#) - Always call addStoredBlock on blocks which have been shifted from one storage to another
- [HDFS-7980](#) - Incremental BlockReport will dramatically slow down the startup of a namenode
- [HDFS-8305](#) - HDFS INotify: the destination field of RenameOp should always end with the file name
- [YARN-3842](#) - NMProxy should retry on NMNotYetReadyException

- [YARN-3467](#) - Expose allocatedMB, allocatedVCores, and runningContainers metrics on running Applications in RM Web UI
- [YARN-3762](#) - FairScheduler: CME on FSParentQueue#getQueueUserAcInfo
- [YARN-3675](#) - FairScheduler: RM quits when node removal races with continuous scheduling on the same node
- [YARN-3491](#) - PublicLocalizer#addResource is too slow.
- [MAPREDUCE-6387](#) - Serialize the recently added Task#encryptedSpillKey field at the end
- [HBASE-13481](#) - Master should respect master (old) DNS/bind related configurations
- [HBASE-13826](#) - Unable to create table when group acls are appropriately set.
- [HBASE-13729](#) - Old hbase.regionserver.global.memstore.upperLimit and hbase.regionserver.global.memstore.lowerLimit properties are ignored if present
- [HBASE-13789](#) - ForeignException should not be sent to the client
- [HBASE-13779](#) - Calling table.exists() before table.get() end up with an empty Result
- [HBASE-13780](#) - Default to 700 for HDFS root dir permissions for secure deployments
- [HBASE-13768](#) - ZooKeeper znodes are bootstrapped with insecure ACLs in a secure configuration
- [HBASE-13767](#) - Allow ZKAcReset to set and not just clear ZK ACLs
- [HBASE-13086](#) - Show ZK root node on Master WebUI
- [HBASE-13413](#) - Create an integration test for Replication
- [HBASE-13611](#) - update clover to work for current versions
- [HIVE-10841](#) - [WHERE col is not null] does not work sometimes for queries with many JOIN statements
- [HIVE-10956](#) - HS2 leaks HMS connections
- [HIVE-10571](#) - HiveMetaStoreClient should close existing thrift connection before its reconnect
- [HIVE-10835](#) - Concurrency issues in JDBC driver
- [HIVE-10802](#) - Table join query with some constant field in select fails
- [HIVE-10538](#) - Fix NPE in FileSinkOperator from hashcode mismatch
- [HIVE-10771](#) - "separatorChar" has no effect in "CREATE TABLE AS SELECT" statement
- [HIVE-10732](#) - Hive JDBC driver does not close operation for metadata queries
- [HIVE-10151](#) - insert into A select from B is broken when both A and B are Acid tables and bucketed the same way
- [HIVE-10483](#) - insert overwrite partition deadlocks on itself with DbTxnManager
- [HIVE-10050](#) - Support overriding memory configuration for AM launched for TempletonControllerJob
- [HIVE-10242](#) - ACID: insert overwrite prevents create table command
- [HIVE-10481](#) - ACID table update finishes but values not really updated if column names are not all lower case
- [HIVE-10150](#) - delete from acidTbl where a in(select a from nonAcidOrcTbl) fails
- [HIVE-10721](#) - SparkSessionManagerImpl leaks SparkSessions [Spark Branch]
- [HIVE-10671](#) - yarn-cluster mode offers a degraded performance from yarn-client [Spark Branch]
- [HIVE-10453](#) - HS2 leaking open file descriptors when using UDFs
- [HIVE-2573](#) - Create per-session function registry
- [HIVE-9520](#) - Create NEXT_DAY UDF
- [HIVE-9143](#) - select user(), current_user()
- [HIVE-5472](#) - support a simple scalar which returns the current timestamp
- [HIVE-10646](#) - ColumnValue does not handle NULL_TYPE
- [HUE-2784](#) - [oozie] Coordinator editor generate wrong Monday cron expression
- [HUE-2793](#) - [JB] Fix Mapper & Reducer counts in job page
- [HUE-2776](#) - [jb] Fix "View All Tasks" pagination
- [HUE-2778](#) - [jobbrowser] Fix "Text Filter" search box in JB "View All Tasks" page
- [HUE-2767](#) - [impala] Issue showing sample data for a table
- [HUE-2754](#) - [oozie] Sqoop action with variable adds an empty argument
- [HUE-2743](#) - [search] Error HTML style leaks in the UI
- [HUE-2587](#) - [jb] Kill jobs in accepted state
- [HUE-2731](#) - [core] Validate that Hue is running in collect data script
- [HUE-2687](#) - [core] Create script to gather hue process info for troubleshooting

- [HUE-2656](#) - [tools] Add cron scripts for restart when mem usage is high
- [HUE-2701](#) - [oozie] Java action relative jar path results in error on submit
- [HUE-2703](#) - [sentry] Make more obvious why a user is not a Sentry admin
- [HUE-2741](#) - [home] Hide the document move dialog
- [HUE-2739](#) - [metastore] Autocomplete with databases/tables with built in names fails
- [HUE-2732](#) - Hue isn't correctly doing add_column migrations with non-blank defaults
- [IMPALA-1963](#): Impala Timestamp ISO-8601 Support.
- [IMPALA-2043](#): skip metadata/testddl.py#test_create_alter_bulk_partition on S3
- [IMPALA-1968](#): Part 1: Improve planner numNodes estimate for remote scans
- [IMPALA-1730](#): reduce scanner thread spinning windows
- [IMPALA-2002](#): Provide way to cache ext data source classes
- [IMPALA-2008](#): Fix wrong warning when insert overwrite to empty table
- [IMPALA-1381](#): Expand set of supported timezones.
- [IMPALA-1952](#): Expand parsing of decimals to include scientific notation
- [SENTRY-227](#) - Fix for "Unsupported entity type DUMMYPARTITION"
- [SQLR-7503](#) - Recovery after ZK session expiration happens in a single thread for all cores in a node
- [SPARK-6299](#) - ClassNotFoundException in standalone mode when running groupByKey with class defined in REPL.
- [SPARK-5522](#) - Accelerate the History Server start

Published Known Issues Fixed

Migrations to MySQL fail if multiple Hue users have the same name but different upper/lower case letters

Bug: CDH-24213

Workaround: None.

Issues Fixed in CDH 5.4.2

Upgrades to CDH 5.4.1 from Releases Earlier than 5.4.0 May Fail

Problem: Because of a change in the implementation of the NameNode metadata upgrade mechanism, upgrading to CDH 5.4.1 from a version lower than 5.4.0 can take an inordinately long time. In a cluster with NameNode high availability (HA) configured and a large number of edit logs, the upgrade can fail, with errors indicating a timeout in the pre-upgrade step on JournalNodes.

What to do:

To avoid the problem: Do not upgrade to CDH 5.4.1; upgrade to CDH 5.4.2 instead.

If you experience the problem: If you have already started an upgrade and seen it fail, contact Cloudera Support. This problem involves no risk of data loss, and manual recovery is possible.

If you have already completed an upgrade to CDH 5.4.1, or are installing a new cluster: In this case you are not affected and can continue to run CDH 5.4.1.

Issues Fixed in CDH 5.4.1

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.4.1:

- [HADOOP-11891](#) - OsSecureRandom should lazily fill its reservoir to avoid open too many file descriptors.
- [HADOOP-11802](#) - DomainSocketWatcher thread terminates sometimes after there is an I/O error during requestShortCircuitShm
- [HADOOP-11724](#) - DistCp throws NPE when the target directory is root.
- [HDFS-7645](#) - Rolling upgrade is restoring blocks from trash multiple times which could cause significant and unnecessary block churn.
- [HDFS-7869](#) - Inconsistency in the return information while performing rolling upgrade
- [HDFS-8127](#) - NameNode Failover during HA upgrade can cause DataNode to finalize upgrade
- [HDFS-3443](#) - Fix NPE when NameNode transition to active during startup.

- [HDFS-7312](#) - Update DistCp v1 to optionally not use tmp location
- [HDFS-8292](#) - Move conditional in fmt_time from dfs-dust.js to status.html
- [HDFS-6673](#) - Add delimited format support to PB OIV tool
- [HDFS-8214](#) - Secondary NN Web UI shows wrong date for Last Checkpoint
- [HDFS-7884](#) - Fix NullPointerException in BlockSender when the generation stamp provided by the client is larger than the one stored in the DataNode
- [HDFS-4448](#) - Allow HA NN to start in secure mode with wildcard address configured
- [HDFS-8070](#) - Fix issue that Pre-HDFS-7915 DFSClient cannot use short circuit on post-HDFS-7915 DataNode
- [HDFS-7915](#) - The DataNode can sometimes allocate a ShortCircuitShm slot and fail to tell the DFSClient about it because of a network error
- [HDFS-7931](#) - DistributedFileSystem should not look for keyProvider in cache if Encryption is disabled
- [HDFS-7916](#) - 'reportBadBlocks' from DataNodes to standby Node BPSERVICEActor goes for infinite loop
- [HDFS-8099](#) - Change "DFSInputStream has been closed already" message to debug log level
- [HDFS-7996](#) - After swapping a volume, BlockReceiver reports ReplicaNotFoundException
- [HDFS-7587](#) - Edit log corruption can happen if append fails with a quota violation
- [HDFS-7881](#) - TestHftpFileSystem#testSeek fails
- [HDFS-7929](#) - inotify is unable to fetch pre-upgrade edit log segments once upgrade starts
- [YARN-3363](#) - Add localization and container launch time to ContainerMetrics at NM to show these timing information for each active container.
- [YARN-3485](#) - FairScheduler headroom calculation doesn't consider maxResources for Fifo and FairShare policies
- [YARN-3464](#) - Race condition in LocalizerRunner kills localizer before localizing all resources
- [YARN-3516](#) - Killing ContainerLocalizer action doesn't take effect when private localizer receives FETCH_FAILURE status.
- [YARN-3021](#) - YARN's delegation-token handling disallows certain trust setups to operate properly over DistCp
- [YARN-3241](#) - FairScheduler handles invalid queue names inconsistently.
- [YARN-2868](#) - FairScheduler: Add a metric for measuring latency of allocating first container for an application
- [YARN-3428](#) - Add debug logs to capture the resources being localized for a container.
- [MAPREDUCE-6339](#) - Job history file is not flushed correctly because isTimerActive flag is not set true when flushTimerTask is scheduled.
- [MAPREDUCE-5710](#) - Running distcp with -delete incurs avoidable penalties
- [MAPREDUCE-6343](#) - JobConf.parseMaximumHeapSizeMB() fails to parse value greater than 2GB expressed in bytes
- [MAPREDUCE-6238](#) - MR2 can't run local jobs with -libjars command options which is a regression from MR1
- [MAPREDUCE-6076](#) - Zero map split input length combined with none zero map split input length may cause MR1 job hung sometimes.
- [HBASE-13374](#) - Small scanners (with particular configurations) do not return all rows
- [HBASE-13269](#) - Limit result array pre-allocation to avoid OOME with large scan caching values
- [HBASE-13335](#) - Update ClientSmallScanner and ClientSmallReversedScanner to use serverHasMoreResults context
- [HBASE-13534](#) - Change HBase master WebUI to explicitly mention if it is a backup master
- [HBASE-13111](#) - truncate_preserve command is failing with undefined method error
- [HBASE-13430](#) - HFiles that are in use by a table cloned from a snapshot may be deleted when that snapshot is deleted
- [HBASE-13546](#) - NPE on RegionServer status page if all masters are down
- [HBASE-13350](#) - Add a debug-warning if we fail HTD checks even if table.sanity.checks is disabled
- [HBASE-13262](#) - ResultScanner doesn't return all rows in Scan
- [HIVE-10452](#) - Avoid sending Beeline prompt+query to the standard output/error only when in script mode.
- [HIVE-10541](#) - Beeline requires newline at the end of each query in a file
- [HIVE-9625](#) - Delegation tokens for HMS are not renewed
- [HIVE-10499](#) - Ensure Session/ZooKeeperClient instances are closed
- [HIVE-10312](#) - SASL.QOP in JDBC URL is ignored for Delegation token Authentication
- [HIVE-10324](#) - Hive metatool should take table_param_key to allow for changes to avro serde's schema url key

- [HIVE-10202](#) - Beeline outputs prompt+query on standard output when used in non-interactive mode
- [HIVE-10087](#) - Beeline's --silent option should suppress query from being echoed when running with -f option
- [HIVE-10098](#) - HS2 local task for map join fails in KMS encrypted cluster
- [HIVE-10146](#) - Add option to not count session as idle if query is running
- [HIVE-10108](#) - Index#getIndexTableName() should return db.index_table_name instead of qualified table name
- [HIVE-10093](#) - Unnecessary HMSHandler initialization for default MemoryTokenStore on HS2
- [HIVE-10085](#) - Lateral view on top of a view throws RuntimeException
- [HIVE-10086](#) - Parquet file using column index access throws error in Hive
- [HIVE-9839](#) - HiveServer2 leaks OperationHandle on async queries which fail at compile phase
- [HIVE-9920](#) - DROP DATABASE IF EXISTS throws exception if database does not exist
- [HIVE-10476](#) - Hive query should fail when it fails to initialize a session in SetSparkReducerParallelism
- [HIVE-10434](#) - Cancel connection when remote Spark driver process has failed
- [HIVE-10473](#) - Spark client is recreated even spark configuration is not changed
- [HIVE-10291](#) - Hive on Spark job configuration needs to be logged
- [HIVE-10143](#) - HS2 fails to clean up Spark client state on timeout
- [HIVE-10073](#) - Runtime exception when querying HBase with Spark
- [HUE-2723](#) - [hive] Listing table information in non default DB fails
- [HUE-2722](#) - [hive] Query returns wrong number of rows when HiveServer2 returns data not encoded properly
- [HUE-2713](#) - [oozie] Deleting a Fork of Fork can break the workflow
- [HUE-2717](#) - [oozie] Coordinator editor does not save non-default schedules
- [HUE-2716](#) - [pig] Scripts fail on hcat auth with org.apache.hive.hcatalog.pig.HCatLoader()
- [HUE-2707](#) - [hive] Allow sample of data on partitioned tables in strict mode
- [HUE-2720](#) - [oozie] Intermittent 500s when trying to view oozie workflow history v1
- [HUE-2712](#) - [oozie] Creating a fork can error
- [HUE-2710](#) - [search] Heatmap select on yelp example errors
- [HUE-2686](#) - [impala] Explain button is erroring
- [HUE-2671](#) - [core] sync_groups_on_login doesn't work with NT Domain
- [IMPALA-1519/IMPALA-1946](#) - Fix wrapping of exprs via a TupleIsNullPredicate with analytics.
- [IMPALA-1900](#) - Assign predicates below analytic functions with a compatible partition by clause for partition pruning.
- [IMPALA-1919](#) - When out_batch->AtCapacity(), avoid calling ProcessBatch in right joins.
- [IMPALA-1960](#) - Illegal reference to non-materialized tuple when query has an empty select-project-join block.
- [IMPALA-1969](#) - OpenSSL init must not be called concurrently.
- [IMPALA-1973](#) - Fixing crash when uninitialized, empty row is added in HdfsTextScanner due to missing newline at the end of file.
- [OOZIE-2218](#) - META-INF directories in the war file have 777 permissions
- [OOZIE-2170](#) - Oozie should automatically set configs to make Spark jobs show up in the Spark History Server
- [SENTRY-699](#) - Memory leak when running Sentry with HiveServer2
- [SENTRY-703](#) - Calls to add_partition fail when passed a Partition object with a null location
- [SENTRY-696](#) - Improve Metastoreplugin Cache Initialization time
- [SENTRY-683](#) - HDFS service client should ensure the kerberos ticket is valid before new service connection
- [SOLR-7478](#) - UpdateLog#close shuts down it's executor with interrupts before running close, possibly preventing a clean close.
- [SOLR-7437](#) - Make HDFS transaction log replication factor configurable.
- [SOLR-7338/SOLR-6583](#) - A reloaded core will never register itself as active after a ZK session expiration.
- [SPARK-7281](#) - No option for AM native library path in yarn-client mode.
- [SPARK-6087](#) - Provide actionable exception if Kryo buffer is not large enough
- [SPARK-6868](#) - Container link broken on Spark UI Executors page when YARN is set to HTTPS_ONLY
- [SPARK-6506](#) - python support in yarn cluster mode requires SPARK_HOME to be set
- [SPARK-6650](#) - ExecutorAllocationManager never stops

- [SPARK-6578](#) - Outbound channel in network library is not thread-safe, can lead to fetch failures
- [SQOOP-2343](#) - AsyncSqlRecordWriter stuck if any exception is thrown out in its close method
- [SQOOP-2286](#) - Ensure Sqoop generates valid avro column names
- [SQOOP-2283](#) - Support usage of --exec and --password-alias
- [SQOOP-2281](#) - Set overwrite on kite dataset
- [SQOOP-2282](#) - Add validation check for --hive-import and --append
- [SQOOP-2257](#) - Import Parquet data into a hive table with --hive-overwrite option does not work
- [ZOOKEEPER-2146](#) - BinaryInputArchive readString should check length before allocating memory
- [ZOOKEEPER-2149](#) - Log client address when socket connection established

Published Known Issues Fixed

As a result of the above fixes, the following issues, previously published as [Known Issues in CDH 5](#) on page 1944, are also fixed.

Apache Hadoop

NameNode cannot use wildcard address in a secure cluster

In a secure cluster, you cannot use a wildcard for the NameNode's RPC or HTTP bind address. For example, `dfs.namenode.http-address` must be a real, routable address and port, not `0.0.0.0.<port>`. This should affect you only if you are running a secure cluster *and* your NameNode needs to bind to multiple local addresses.

Bug: [HDFS-4448](#)

Cloudera Bug: CDH-9991

Workaround: None

Offline Image Viewer (OIV) tool regression: missing Delimited outputs.

Bugs: [HDFS-6673](#), [HDFS-5952](#)

Cloudera Bug: CDH-20259

Severity: Medium

Workaround: Set up `dfs.namenode.legacy-oiv-image.dir` to an appropriate directory on the secondary NameNode (or standby NameNode in an HA configuration), and use `hdfs oiv_legacy` to process the legacy format of the OIV `fsimage`.

Apache HBase

Setting `maxResultSize` Incorrectly On a Scan May Cause Client Data Loss

Scanners may not return all the results from a region if a scan is configured with a `maxResultSize` limit that could be reached before the caching limit. Results are missed because the scanner jumps to the next region preemptively.

The default value for `maxResultSize` is `Long.MAX_VALUE` and the default value of caching is 100, so with the default configuration, the caching limit will always be reached before the `maxResultSize` and the issue will not appear. If the `maxResultSize` is configured to any limit that may be reached before the caching limit, the issue may occur.

Bug: [HBASE-13262](#)

Severity: Low

Workaround: Never configure a scan with a `maxResultSize` other than `Long.MAX_VALUE` (never change it from its default value) because that will ensure that the `maxResultSize` limit is never reached before the caching limit.

Apache Hive

Hive metatool does not fix Avro schema URL setting in an HDFS HA upgrade

When you upgrade Hive in an HDFS HA configuration, and the `avro.schema.url` is set in an Avro table's properties instead of the SerDe properties, the metatool will not correct the problem.

Bug: [HIVE-10324](#)

Cloudera Bug: CDH-26976

Workaround: Use `alter table.. set tblproperties` to fix the `avro.schema.url`.

Hive metastore `getIndexTableName` returns qualified table name

In CDH 5.4.0, `getIndexTableName` returns a qualified table name such as

```
database_name
.
index_table_name
```

whereas in previous releases it returns an unqualified table name, such as

```
index_table_name
```

Bug: [HIVE-10108](#)

Cloudera Bug: CDH-26508

Workaround: None

HiveServer2 has an unexpected Derby metastore directory in secure clusters

Bug: [HIVE-10093](#)

Cloudera Bug: CDH-26453

Workaround: None. Ignore the Derby database.

Apache Oozie

Spark jobs run from the Spark action don't show up in the Spark History Server or properly link to it from the Spark AM

Bug: [OOZIE-2170](#)

Cloudera Bug: CDH-25513

Severity: Low

Workaround: Specify these configuration properties in the `spark-opts` element of your Spark action in the `workflow.xml` file:

```
--conf spark.yarn.historyServer.address=http://SPH:18088 --conf
spark.eventLog.dir=hdfs://NN:8020/user/spark/applicationHistory --conf
spark.eventLog.enabled=true
```

where *SPH* is the hostname of the Spark History Server and *NN* is the hostname of the NameNode. You can also find these values in `/etc/spark/conf/spark-defaults.conf` on the gateway host when Spark is installed from Cloudera Manager.

Apache Sentry

Hive binding should support enforcing URI privilege for transforms

Bug: [SENTRY-598](#)

Severity: Medium

Workaround: None.

Issues Fixed in CDH 5.4.0

The following topics describe known issues fixed in CDH 5.4.0.

For the latest Impala fixed issues, see [Issues Fixed in CDH 5.4 / Impala 2.2](#) on page 2436.

Apache Hadoop

HDFS

After upgrade from a release earlier than CDH 5.2.0, storage IDs may no longer be unique

As of CDH 5.2, each storage volume on a DataNode should have its own unique `storageID`, but in clusters upgraded from CDH 4, or CDH 5 releases earlier than CDH 5.2.0, each volume on a given DataNode shares the same `storageID`, because the HDFS upgrade does not properly update the IDs to reflect the new naming scheme. This causes problems with load balancing. The problem affects only clusters upgraded from CDH 5.1.x and earlier to CDH 5.2 or later. Clusters that are new as of CDH 5.2.0 or later do not have the problem.

Bug: [HDFS-7575](#)

Cloudera Bug: CDH-24155

Severity: Medium

Workaround: Upgrade to a later or patched version of CDH.

Apache Hive

UDF infile() does not accept arguments of type CHAR or VARCHAR

Bug: [HIVE-6637](#)

Severity: Low

Workaround: Cast the argument to type String.

Hive's Decimal type cannot be stored in Parquet and Avro

Tables containing decimal columns cannot use Parquet or the Avro storage engine.

Bug: [HIVE-6367](#) and [HIVE-5823](#)

Severity: Low

Workaround: Use a different file format.

Apache Oozie

Executing oozie job -config properties file -dryrun fails because of a code defect in argument parsing

Bug: [OOZIE-1878](#)

Cloudera Bug: CDH-25537

Severity: Low

Workaround: None.

When you use Hive Server 2 from Oozie, Oozie won't collect or print out the Hadoop Job IDs of any jobs launched by Hive Server 2

Cloudera Bug: CDH-12777, CDH-12891

Severity: Low

Workaround: You can get the Hadoop IDs from the Resource Manager or JobTracker.

Cloudera Search

Spark indexer failed if configured to use security.

Spark indexing jobs failed when Kerberos authentication was enabled.

With Search for CDH 5.4 and later, Spark indexing jobs succeed, even when Kerberos authentication is required.

Bug: None.

Severity: Medium.

Workaround: Disable Kerberos authentication or use another indexer.

Mapper-only HBase batch indexer failed if configured to use security.

Attempts to complete an HBase batch indexing job failed when Kerberos authentication was enabled and reducers were set to 0.

With Search for CDH 5.4 and later, mapper-only HBase batch indexer succeeds, even when Kerberos authentication is required.

Bug: None.

Severity: Medium.

Workaround: Either disable Kerberos authentication or use one or more reducers.

Shard splitting support is experimental.

Cloudera anticipated shard splitting to function as expected with Cloudera Search, but this interaction had not been thoroughly tested.

As of the release of Search for CDH 5.4, additional testing of shard splitting has been completed, so this functionality can be safely used.

Cloudera Bug: CDH-11024

Severity: Low

Workaround: Use shard splitting for test and development purposes, but be aware of the risks of using shard splitting in production environments. To avoid using shard splitting, use the source data to create a new index with a new sharding count by re-indexing the data to a new collection. You can enable this using the MapReduceIndexerTool.

TrieDateField defaulted OMIT_NORMS to True.

All primitive field types were intended to omit norms by default with schema version 1.5 or higher. This change was not applied to `TrieDateField`.

With Search for CDH 5.4, `TrieDateField` is set to omit norms by default.

Bug: SOLR-6211

Severity: Low

Fields or Types outside <field> or <types> tags are silently ignored.

In previous releases, Solr silently ignored definitions such as `<fieldType>`, `<field>`, and `<copyField>` if those definitions were not contained in `<fields>` or `<types>` tags.

With Search 5.4 for CDH, these tags are no longer required for definitions to be included. These tags are supported so either style may be implemented.

Bug: SOLR-5228

Apache Sentry (incubating)

INSERT OVERWRITE LOCAL fails if you use only the Linux pathname

Cloudera Bug: CDH-13732

Severity: Low

Workaround: Prefix the path of the local file with `file://` when using `INSERT OVERWRITE LOCAL`.

INSERT OVERWRITE and CREATE EXTERNAL commands fail because of HDFS URI permissions

When you use Sentry to secure Hive, and use HDFS URIs in a HiveQL statement, the query will fail with an HDFS permissions error unless you specify the NameNode and port.

Cloudera Bug: CDH-13728

Severity: Low

Workaround: Specify the NameNode and port, where applicable, in the URI; for example specify `hdfs://nn-uri:port/user/warehouse/hive/tab` rather than simply `/user/warehouse/hive/tab`. In a high-availability deployment, specify the value of `FS.defaultFS`.

Issues Fixed in CDH 5.3.x

The following topics describe issues fixed in CDH 5.3.x, from newest to oldest release. You can also review [What's New In CDH 5.3.x](#) on page 1845 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.3.10

CDH 5.3.10 fixes the following issues.

Apache Hadoop

FSImage may get corrupted after deleting snapshot

Bug: [HDFS-9406](#)

Cloudera Bug: CDH-33224

When deleting a snapshot that contains the last record of a given INode, the fsmage may become corrupt because the create list of the snapshot diff in the previous snapshot and the child list of the parent INodeDirectory are not cleaned.

Apache HBase

The ReplicationCleaner process can abort if its connection to ZooKeeper is inconsistent

Bug: [HBASE-15234](#)

If the connection with ZooKeeper is inconsistent, the ReplicationCleaner may abort, and the following event is logged by the HMaster:

```
WARN org.apache.hadoop.hbase.replication.master.ReplicationLogCleaner: Aborting
ReplicationLogCleaner
because Failed to get list of replicators
```

Unprocessed WALs accumulate.

The `seekBefore()` method calculates the size of the previous data block by assuming that data blocks are contiguous, and HFile v2 and higher store Bloom blocks and leaf-level INode blocks with the data. As a result, reverse scans do not work when Bloom blocks or leaf-level INode blocks are present when HFile v2 or higher is used.

Workaround: Restart the HMaster occasionally. The ReplicationCleaner restarts if necessary and process the unprocessed WALs.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.3.10:

- [HADOOP-7713](#) - dfs -count -q should label output column
- [HADOOP-8944](#) - Shell command fs -count should include human readable option
- [HADOOP-10406](#) - TestIPC.testIpcWithReaderQueuing may fail
- [HADOOP-12200](#) - TestCryptoStreamsWithOpensslAesCtrCryptoCodec should be skipped in non-native profile
- [HADOOP-12240](#) - Fix tests requiring native library to be skipped in non-native profile
- [HADOOP-12280](#) - Skip unit tests based on maven profile rather than NativeCodeLoader.isNativeCodeLoaded
- [HADOOP-12418](#) - TestRPC.testRPCInterruptedSimple fails intermittently
- [HADOOP-12464](#) - Interrupted client may try to fail over and retry
- [HADOOP-12468](#) - Partial group resolution failure should not result in user lockout
- [HADOOP-12559](#) - KMS connection failures should trigger TGT renewal
- [HADOOP-12604](#) - Exception may be swallowed in KMSClientProvider
- [HADOOP-12605](#) - Fix intermittent failure of TestIPC.testIpcWithReaderQueuing
- [HADOOP-12682](#) - Fix TestKMS#testKMSRestart* failure
- [HADOOP-12699](#) - TestKMS#testKMSPProvider intermittently fails during 'test rollover draining'
- [HADOOP-12715](#) - TestValueQueue#testgetAtMostPolicyALL fails intermittently
- [HADOOP-12736](#) - TestTimedOutTestsListener#testThreadDumpAndDeadlocks sometimes times out
- [HADOOP-12788](#) - OpensslAesCtrCryptoCodec should log which random number generator is used

- [HDFS-6533](#) - TestBPOfferService#testBasicFunctionalitytest fails intermittently
- [HDFS-6673](#) - Add delimited format support to PB OIV tool
- [HDFS-6799](#) - The invalidate method in SimulatedFSDataset failed to remove (invalidate) blocks from the file system
- [HDFS-7423](#) - Various typos and message formatting fixes in nfs daemon and doc
- [HDFS-7553](#) - Fix the TestDFSUpgradeWithHA due to BindException
- [HDFS-7990](#) - IBR delete ack should not be delayed
- [HDFS-8211](#) - DataNode UUID is always null in the JMX counter
- [HDFS-8646](#) - Prune cached replicas from DatanodeDescriptor state on replica invalidation
- [HDFS-9092](#) - NFS silently drops overlapping write requests and causes data copying to fail
- [HDFS-9250](#) - Add Precondition check to LocatedBlock#addCachedLoc
- [HDFS-9347](#) - Invariant assumption in TestQuorumJournalManager.shutdown() is wrong
- [HDFS-9358](#) - TestNodeCount#testNodeCount timed out
- [HDFS-9364](#) - Unnecessary DNS resolution attempts when creating NameNodeProxies
- [HDFS-9406](#) - FSImage may get corrupted after deleting snapshot
- [HDFS-9949](#) - Add a test case to ensure that the DataNode does not regenerate its UUID when a storage directory is cleared
- [MAPREDUCE-6302](#) - Incorrect headroom can lead to a deadlock between map and reduce allocations
- [MAPREDUCE-6387](#) - Serialize the recently added Task#encryptedSpillKey field at the end
- [MAPREDUCE-6460](#) - TestRMContainerAllocator.testAttemptNotFoundCausesRMCommunicatorException fails
- [YARN-2377](#) - Localization exception stack traces are not passed as diagnostic info
- [YARN-2785](#) - Fixed intermittent TestContainerResourceUsage failure
- [YARN-3024](#) - LocalizerRunner should give DIE action when all resources are localized
- [YARN-3074](#) - Nodemanager dies when localizer runner tries to write to a full disk
- [YARN-3464](#) - Race condition in LocalizerRunner kills localizer before localizing all resources.
- [YARN-3516](#) - Killing ContainerLocalizer action does not take effect when private localizer receives FETCH_FAILURE status
- [YARN-3727](#) - For better error recovery, check if the directory exists before using it for localization
- [YARN-3762](#) - FairScheduler: CME on FSParentQueue#getQueueUserAclInfo
- [YARN-4204](#) - ConcurrentModificationException in FairSchedulerQueueInfo
- [YARN-4235](#) - FairScheduler PrimaryGroup does not handle empty groups returned for a user
- [YARN-4354](#) - Public resource localization fails with NPE
- [YARN-4380](#) - TestResourceLocalizationService.testDownloadingResourcesOnContainerKill fails intermittently
- [YARN-4393](#) - Fix intermittent test failure for TestResourceLocalizationService#testFailedDirsResourceRelease
- [YARN-4613](#) - Fix test failure in TestClientRMService#testGetClusterNodes
- [YARN-4717](#) - TestResourceLocalizationService.testPublicResourceInitializesLocalDir fails Intermittently due to IllegalArgumentException from cleanup
- [HBASE-10153](#) - Improve VerifyReplication to compute BADROWS more accurately
- [HBASE-11394](#) - AmendReplication can have data loss if peer id contains hyphen
- [HBASE-11394](#) - Replication can have data loss if peer id contains hyphen "-"
- [HBASE-11992](#) - Backport HBASE-11367 (Pluggable replication endpoint) to 0.98
- [HBASE-12136](#) - Race condition between client adding tableCF replication znode and server triggering TableCFsTracker
- [HBASE-12150](#) - Backport replication changes from HBASE-12145
- [HBASE-12336](#) - RegionServer failed to shutdown for NodeFailoverWorker thread
- [HBASE-12631](#) - Backport HBASE-12576 (Add metrics for rolling the HLog if there are too few DNs in the write pipeline) to 0.98
- [HBASE-12658](#) - Backport HBASE-12574 (Update replication metrics to not do so many map look ups) to 0.98
- [HBASE-12865](#) - WALs may be deleted before they are replicated to peers
- [HBASE-13035](#) - Backport HBASE-12867 Shell does not support custom replication endpoint specification
- [HBASE-13084](#) - Add labels to VisibilityLabelsCache asynchronously causes TestShell flakey
- [HBASE-13437](#) - ThriftServer leaks ZooKeeper connections

- [HBASE-13703](#) - ReplicateContext should not be a member of ReplicationSource
- [HBASE-13746](#) - list_replicated_tables command is not listing table in hbase shell
- [HBASE-14146](#) - Fix Once replication sees an error it slows down forever
- [HBASE-14501](#) - NPE in replication with TDE
- [HBASE-14621](#) - ReplicationLogCleaner gets stuck when a regionserver crashes
- [HBASE-14923](#) - VerifyReplication should not mask the exception during result comparison
- [HBASE-15019](#) - Replication stuck when HDFS is restarted
- [HBASE-15032](#) - hbase shell scan filter string assumes UTF-8 encoding
- [HBASE-15035](#) - bulkloading hfiles with tags that require splits do not preserve tags
- [HBASE-15052](#) - Use EnvironmentEdgeManager in ReplicationSource
- [HIVE-7524](#) - Enable auto conversion of SMBjoin in presence of constant propagate optimization
- [HIVE-7575](#) - GetTables thrift call is very slow
- [HIVE-8115](#) - Fixing text failures caused in CDH
- [HIVE-8115](#) - Hive select query hang when fields contain map
- [HIVE-8184](#) - Inconsistency between colList and columnExprMap when ConstantPropagate is applied to subquery
- [HIVE-9112](#) - Query may generate different results depending on the number of reducers
- [HIVE-9500](#) - Support nested structs over 24 levels
- [HIVE-9860](#) - MapredLocalTask/SecureCmdDoAs leaks local files
- [HIVE-10956](#) - Fallout fix from backport to CDH 5.3.x
- [HIVE-11977](#) - Hive should handle an external avro table with zero length files present
- [HIVE-12388](#) - GetTables cannot get external tables when TABLE type argument is given
- [HIVE-12406](#) - HIVE-9500 introduced incompatible change to LazySimpleSerDe public interface
- [HIVE-12713](#) - Miscellaneous improvements in driver compile and execute logging
- [HIVE-12790](#) - Metastore connection leaks in HiveServer2
- [HIVE-12946](#) - alter table should also add default scheme and authority for the location similar to create table
- [HUE-2767](#) - [impala] Issue showing sample data for a table
- [HUE-2941](#) - [hadoop] Cache the active RM HA
- [IMPALA-1702](#) - "invalidate metadata" can cause duplicate TableIds (issue not entirely fixed, but now fails gracefully)
- [IMPALA-2125](#) - Improve perf when reading timestamps from parquet files written by hive
- [IMPALA-2565](#) - Planner tests are flaky due to file size mismatches
- [IMPALA-3095](#) - Allow additional Kerberos users to be authorized to access internal APIs
- [OOZIE-2432](#) - TestPurgeXCommand fails
- [SENTRY-565](#) - Improve performance of filtering Hive SHOW commands
- [SENTRY-780](#) - HDFS Plugin should not execute path callbacks for views
- [SENTRY-835](#) - Drop table leaves a connection open when using metastorelistener
- [SENTRY-885](#) - DB name should be case insensitive in HDFS sync plugin.
- [SENTRY-936](#) - getGroup and getUser should always return original hdfs values for paths in prefix which are not sentry managed
- [SENTRY-944](#) - Setting HDFS rules on Sentry managed hdfs paths should not affect original hdfs rules
- [SENTRY-957](#) - Exceptions in MetastoreCacheInitializer should probably not prevent HMS from starting up
- [SENTRY-988](#) - It is better to let SentryAuthorization setter path always fall through and update HDFS
- [SENTRY-994](#) - SentryAuthorizationInfoX should override isSentryManaged
- [SENTRY-1002](#) - PathsUpdate.parsePath(path) will throw an NPE when parsing relative paths
- [SENTRY-1044](#) - Tables with non-hdfs locations break HMS startup
- [SPARK-12617](#) - [PYSPARK] Move Py4jCallbackConnectionCleaner to Streaming

Issues Fixed in CDH 5.3.9

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera ("Cloudera Products"), including core Apache

Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled “Arbitrary remote code execution with InvokerTransformer” and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.3.9:

- [FLUME-2841](#) - Upgrade commons-collections to 3.2.2
- [HADOOP-12577](#) - Bumped up commons-collections version to 3.2.2 to address a security flaw
- [HDFS-7785](#) - Improve diagnostics information for HttpPutFailedException
- [HDFS-7798](#) - Checkpointing failure caused by shared KerberosAuthenticator
- [HDFS-7871](#) - NameNodeEditLogRoller can keep printing 'Swallowing exception' message
- [HDFS-9123](#) - Copying from the root to a subdirectory should be forbidden
- [HDFS-9273](#) - ACLs on root directory may be lost after NN restart
- [HDFS-9332](#) - Fix Precondition failures from NameNodeEditLogRoller while saving namespace
- [HDFS-9470](#) - Encryption zone on root not loaded from fsimage after NN restart
- [MAPREDUCE-6191](#) - Improve clearing stale state of Java serialization testcase
- [MAPREDUCE-6233](#) - org.apache.hadoop.mapreduce.TestLargeSort.testLargeSort failed in trunk
- [MAPREDUCE-6549](#) - Multibyte delimiters with LineRecordReader cause duplicate records
- [YARN-3564](#) - Fix TestContainerAllocation.testAMContainerAllocationWhenDNSUnavailable fails randomly
- [YARN-3602](#) - TestResourceLocalizationService.testPublicResourceInitializesLocalDir fails Intermittently due to IOException from cleanup
- [YARN-3675](#) - FairScheduler: RM quits when node removal races with continuous-scheduling on the same node
- [HBASE-13134](#) - mutateRow and checkAndMutate APIs do not throw region level exceptions
- [HBASE-14196](#) - Thrift server idle connection timeout issue
- [HBASE-14283](#) - Reverse scan does not work with HFile inline index/bloom blocks
- [HBASE-14533](#) - Thrift client gets "AsyncProcess: Failed to get region location closed"
- [HBASE-14799](#) - Commons-collections object deserialization remote command execution vulnerability
- [HIVE-6099](#) - Multi insert does not work properly with distinct count
- [HIVE-7146](#) - posexplode() UDTF fails with a NullPointerException on NULL columns
- [HIVE-8612](#) - Support metadata result filter hooks
- [HIVE-9475](#) - HiveMetastoreClient.tableExists does not work
- [HIVE-10895](#) - ObjectStore does not close Query objects in some calls, causing a potential leak in some metastore db resources

- [HIVE-11255](#) - `get_table_objects_by_name()` in `HiveMetaStore.java` needs to retrieve table objects in multiple batches
- [HIVE-12378](#) - Exception on `HBaseSerDe.serialize` binary field
- [HUE-3035](#) - [beeswax] Optimize sample data query for partitioned tables
- [IMPALA-1746](#) - `QueryExecState` does not check for query cancellation or errors
- [IMPALA-1756](#) - Constant filter expressions are not checked for errors and state cleanup not done before throwing exception
- [IMPALA-1917](#) - DCHECK on destroying an `ExprContext`
- [IMPALA-2141](#) - `UnionNode::getNext()` does not check for query errors
- [IMPALA-2264](#) - Fix edge cases for decimal/integer cast
- [IMPALA-2514](#) - DCHECK on destroying an `ExprContext`
- [OOZIE-2413](#) - Kerberos credentials can expire if the KDC is slow to respond
- [PIG-3641](#) - Split "otherwise" producing incorrect output when combined with `ColumnPruning`
- [SPARK-11484](#) - [WEBUI] Using `proxyBase` set by spark instead of env
- [SPARK-11652](#) - [CORE] Remote code execution with `InvokerTransformer`

Issues Fixed in CDH 5.3.8

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.3.8:

- [CRUNCH-525](#) - Correct (more) accurate default scale factors for built-in `MapFn` implementations
- [CRUNCH-527](#) - Use hash smearing for partitioning
- [CRUNCH-528](#) - Improve Pair comparison
- [CRUNCH-535](#) - call `initCredentials` on the job
- [CRUNCH-536](#) - Refactor `CrunchControlledJob.Hook` interface and make it client-accessible
- [CRUNCH-539](#) - Fix reading `WritableComparables` bimap
- [CRUNCH-540](#) - Make `AvroReflectDeepCopier` serializable
- [CRUNCH-542](#) - Eliminate flaky `Scrunch` sampling test.
- [CRUNCH-543](#) - Have `AvroPathPerKeyTarget` handle child directories properly
- [CRUNCH-544](#) - Improve performance/serializability of materialized `toMap`.
- [CRUNCH-547](#) - Properly handle nullability for Avro union types
- [CRUNCH-548](#) - Have the `AvroReflectDeepCopier` use the class of the source object when constructing new instances instead of the target class
- [CRUNCH-551](#) - Make the use of Configuration objects consistent in `CrunchInputSplit` and `CrunchRecordReader`
- [CRUNCH-553](#) - Fix record drop issue that can occur w/`From.formattedFile` TableSources
- [FLUME-1934](#) - Spooling Directory Source dies on encountering zero-byte files.
- [FLUME-2095](#) - JMS source with TIBCO
- [FLUME-2385](#) - Remove incorrect log message at INFO level in Spool Directory Source.
- [FLUME-2753](#) - Error when specifying empty replace string in Search and Replace Interceptor
- [HADOOP-11105](#) - `MetricsSystemImpl` could leak memory in registered callbacks
- [HADOOP-11446](#) - `S3AOutputStream` should use shared thread pool to avoid `OutOfMemoryError`
- [HADOOP-11463](#) - Replace method-local `TransferManager` object with `S3AFileSystem#transfers`.
- [HADOOP-11584](#) - s3a file block size set to 0 in `getFileStatus`.
- [HADOOP-11607](#) - Reduce log spew in `S3AFileSystem`.
- [HADOOP-12317](#) - Applications fail on NM restart on some linux distro because NM container recovery declares AM container as LOST
- [HADOOP-12404](#) - Disable caching for `JarURLConnection` to avoid sharing `JarFile` with other users when loading resource from URL in Configuration class
- [HADOOP-12413](#) - `AccessControlList` should avoid calling `getGroupNames` in `isUserInList` with empty groups
- [HDFS-7978](#) - Add `LOG.isDebugEnabled()` guard for some `LOG.debug(..)`
- [HDFS-8384](#) - Allow NN to startup if there are files having a lease but are not under construction

- [HDFS-8964](#) - When validating the edit log, do not read at or beyond the file offset that is being written
- [HDFS-8965](#) - Harden edit log reading code against out of memory errors
- [MAPREDUCE-5918](#) - LineRecordReader can return the same decompressor to CodecPool multiple times
- [MAPREDUCE-5948](#) - org.apache.hadoop.mapred.LineRecordReader does not handle multibyte record delimiters well
- [MAPREDUCE-6277](#) - Job can post multiple history files if attempt loses connection to the RM
- [MAPREDUCE-6439](#) - AM may fail instead of retrying if RM shuts down during the allocate call.
- [MAPREDUCE-6481](#) - LineRecordReader may give incomplete record and wrong position/key information for uncompressed input sometimes
- [MAPREDUCE-6484](#) - Yarn Client uses local address instead of RM address as token renewer in a secure cluster when RM HA is enabled
- [YARN-3385](#) - Fixed a race-condition in ResourceManager's ZooKeeper based state-store to avoid crashing on duplicate deletes
- [YARN-3469](#) - ZKRMStateStore: Avoid setting watches that are not required.
- [YARN-3990](#) - AsyncDispatcher may overloaded with RMAppNodeUpdateEvent when Node is connected/disconnected
- [HBASE-12639](#) - Backport HBASE-12565 Race condition in HRegion.batchMutate() causes partial data to be written when region closes
- [HBASE-13217](#) - Procedure fails due to ZK issue
- [HBASE-13388](#) - Handling NullPointerException in ZKProcedureMemberRpcs while getting ZNode data
- [HBASE-13437](#) - ThriftServer leaks ZooKeeper connections
- [HBASE-13471](#) - Fix a possible infinite loop in doMiniBatchMutation
- [HBASE-13684](#) - Allow mlockagent to be used when not starting as root
- [HBASE-13885](#) - ZK watches leaks during snapshots.
- [HBASE-14045](#) - Bumping thrift version to 0.9.2.
- [HBASE-14302](#) - TableSnapshotInputFormat should not create back references when restoring snapshot
- [HBASE-14354](#) - Minor improvements for usage of the mlock agent
- [HIVE-4867](#) - Deduplicate columns appearing in both the key list and value list of ReduceSinkOperator
- [HIVE-7012](#) - Wrong RS de-duplication in the ReduceSinkDeDuplication Optimizer
- [HIVE-8162](#) - Dynamic sort optimization propagates additional columns even in the absence of order by
- [HIVE-8398](#) - ExprNodeColumnDesc cannot be cast to ExprNodeConstantDesc
- [HIVE-8404](#) - ColumnPruner doesnt prune columns from limit operator
- [HIVE-8560](#) - SerDes that do not inherit AbstractSerDe do not get table properties during initialize()
- [HIVE-9195](#) - CBO changes constant to column type
- [HIVE-9450](#) - Merge[Parquet] Check all data types work for Parquet in Group
- [HIVE-9613](#) - Left join query plan outputs wrong column when using subquery
- [HIVE-9984](#) - JoinReorder's getOutputSize is exponential
- [HIVE-10319](#) - Hive CLI startup takes a long time with a large number of databases
- [HIVE-10572](#) - Improve Hive service test to check empty string
- [HIVE-11077](#) - partOfExchange partition does not properly populate fields for post/pre execute hooks.
- [HIVE-11172](#) - Retrofit Q-Test + Vectorization wrong results for aggregate query with where clause without group by
- [HIVE-11172](#) - Vectorization wrong results for aggregate query with where clause without group by
- [HIVE-11174](#) - Hive does not treat floating point signed zeros as equal (-0.0 should equal 0.0 according to IEEE floating point spec)
- [HIVE-11203](#) - Beeline force option does not force execution when errors occurred in a script.
- [HIVE-11216](#) - UDF GenericUDFMapKeys throws NPE when a null map value is passed in
- [HIVE-11271](#) - java.lang.IndexOutOfBoundsException when union all with if function
- [HIVE-11288](#) - Avro SerDe InstanceCache returns incorrect schema
- [HIVE-11333](#) - ColumnPruner prunes columns of UnionOperator that should be kept
- [HIVE-11590](#) - AvroDeserializer is very chatty
- [HIVE-11657](#) - HIVE-2573 introduces some issues during metastore init (and CLI init)

- [HIVE-11695](#) - If user have no permission to create LOCAL DIRECTORY `hql` does not throw any exception and fail silently.
- [HIVE-11696](#) - Exception when table-level serde is Parquet while partition-level serde is JSON
- [HIVE-11816](#) - Upgrade groovy to 2.4.4
- [HIVE-11824](#) - Insert to local directory causes staging directory to be copied
- [HIVE-11995](#) - Remove repetitively setting permissions in insert/load overwrite partition
- [HUE-2880](#) - [hadoop] Fix uploading large files to a kerberized HTTPFS
- [HUE-2893](#) - [desktop] Backport CherryPy SSL file upload fix
- [IMPALA-1929](#) - Avoiding a DCHECK of NULL hash table in spilled right joins
- [IMPALA-2133](#) - Properly unescape string value for HBase filters
- [IMPALA-2165](#) - Avoid cardinality 0 in scan nodes of small tables and low selectivity
- [IMPALA-2178](#) - fix Expr::ComputeResultsLayout() logic
- [IMPALA-2314](#) - LargestSpilledPartition was not checking if partition is closed
- [IMPALA-2364](#) - Wrong DCHECK in PHJ::ProcessProbeBatch
- [KITE-1053](#) - Fix int overflow bug in FS writer.
- [KITE-1074](#) - Partial updates aka Atomic updates with loadSolr aren't recognized with Solrcloud
- [MAHOUT-1771](#) - Cluster dumper omits indices and 0 elements for dense vector or sparse containing 0s, this closes apache/mahout#158
- [MAHOUT-1771](#) - Cluster dumper omits indices and 0 elements for dense vector or sparse containing 0s closes apache/mahout #158
- [PIG-4024](#) - TestPigStreamingUDF and TestPigStreaming fail on IBM JDK
- [PIG-4326](#) - AvroStorageSchemaConversionUtilities does not properly convert schema for maps of arrays of records
- [PIG-4338](#) - Fix test failures with JDK8
- [SENTRY-799](#) - unit test forFix TestDbEndToEnd flaky test - drop table/dbs before creating
- [SENTRY-878](#) - collect_list missing from HIVE_UDF_WHITE_LIST
- [SENTRY-893](#) - Synchronize calls in SentryClient and create sentry client once per request in SimpleDBProvider
- [SOLR-5496](#) - Ensure all http CMs get shutdown.
- [SOLR-7956](#) - There are interrupts on shutdown in places that can cause ChannelAlreadyClose
- [SOLR-7999](#) - SolrRequetParserTest#testStreamURL started failing.
- [SPARK-6480](#) - [CORE] histogram() bucket function is wrong in some simple edge cases
- [SPARK-6880](#) - [CORE]Fixed null check when all the dependent stages are cancelled due to previous stage failure
- [SPARK-8606](#) - Prevent exceptions in RDD.getPreferredLocations() from crashing DAGScheduler

Issues Fixed in CDH 5.3.6

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.3.6:

- [CRUNCH-516](#) - Scrunch needs some additional null checks
- [CRUNCH-508](#) - Improve performance of Scala Enumeration counters in Scrunch
- [CRUNCH-514](#) - AvroDerivedDeepCopier should initialize delegate MapFns
- [CRUNCH-530](#) - Fix object reuse bug in GenericRecordToTuple
- [HADOOP-12158](#) - Improve error message in TestCryptoStreamsWithOpensslAesCtrCryptoCodec when OpenSSL is not installed
- [HADOOP-11711](#) - Provide a default value for AES/CTR/NoPadding CryptoCodec classes
- [HADOOP-12103](#) - Small refactoring of DelegationTokenAuthenticationFilter to allow code sharing
- [HADOOP-8151](#) - Error handling in snappy decompressor throws invalid exceptions
- [HADOOP-11969](#) - ThreadLocal initialization in several classes is not thread safe
- [HDFS-7443](#) - Datanode upgrade to BLOCKID_BASED_LAYOUT fails if duplicate block files are present in the same volume
- [HDFS-8337](#) - Accessing httpfs via webhdfs doesn't work from a jar with kerberos
- [HDFS-7546](#) - Document, and set an accepting default for dfs.namenode.kerberos.principal.pattern

- [HDFS-8656](#) - Preserve compatibility of ClientProtocol#rollingUpgrade after finalization
- [HDFS-7894](#) - Rolling upgrade readiness is not updated in jmx until query command is issued.
- [HDFS-8127](#) - NameNode Failover during HA upgrade can cause DataNode to finalize upgrade
- [HDFS-3443](#) - Fix NPE when namenode transition to active during startup by adding checkNNStartup() in NameNodeRpcServer
- [YARN-3143](#) - RM Apps REST API can return NPE or entries missing id and other fields
- [HBASE-13995](#) - ServerName is not fully case insensitive
- [HBASE-13430](#) - HFiles that are in use by a table cloned from a snapshot may be deleted when that snapshot is deleted
- [HBASE-12539](#) - HFileLinkCleaner logs are uselessly noisy
- [HBASE-11898](#) - CoprocessorHost.Environment should cache class loader instance
- [HBASE-13826](#) - Unable to create table when group acls are appropriately set.
- [HBASE-13241](#) - Add tests for group level grants
- [HBASE-13239](#) - HBase grant at specific column level does not work for Groups
- [HBASE-13789](#) - ForeignException should not be sent to the client
- [HBASE-13779](#) - Calling table.exists() before table.get() end up with an empty Result
- [HBASE-13780](#) - Default to 700 for HDFS root dir permissions for secure deployments
- [HBASE-13768](#) - ZooKeeper znodes are bootstrapped with insecure ACLs in a secure configuration
- [HBASE-13767](#) - Allow ZKAcLReset to set and not just clear ZK ACLs
- [HBASE-13086](#) - Show ZK root node on Master WebUI
- [HBASE-13342](#) - Fix incorrect interface annotations
- [HBASE-13162](#) - Add capability for cleaning hbase acls to hbase cleanup script.
- [HBASE-12641](#) - Grant all permissions of hbase zookeeper node to hbase superuser in a secure cluster
- [HBASE-12414](#) - Move HFileLink.exists() to base class
- [HIVE-11150](#) - Remove wrong warning message related to chgrp
- [HIVE-8318](#) - Null Scan optimizer throws exception when no partitions are selected
- [HIVE-7385](#) - Optimize for empty relation scans
- [HIVE-7299](#) - Enable metadata only optimization on Tez
- [HIVE-10808](#) - Inner join on Null throwing Cast Exception
- [HIVE-9087](#) - The move task does not handle properly in the case of loading data from the local file system path.
- [HIVE-9325](#) - Handle the case of insert overwrite statement with a qualified path that the destination path does not have a schema.
- [HIVE-9349](#) - Remove the schema in the getQualifiedPathWithoutSchemeAndAuthority method
- [HIVE-9328](#) - Tests cannot move files due to change on HIVE-9325
- [HIVE-6024](#) - Load data local inpath unnecessarily creates a copy task
- [HIVE-10841](#) - [WHERE col is not null] does not work sometimes for queries with many JOIN statements
- [HIVE-9620](#) - Cannot retrieve column statistics using HMS API if column name contains uppercase characters
- [HIVE-8863](#) - Cannot drop table with uppercase name after "compute statistics for columns"
- [HIVE-10629](#) - Dropping table in an encrypted zone does not drop warehouse directory
- [HIVE-10630](#) - Renaming tables across encryption zones renames table even though the operation throws error
- [HIVE-10956](#) - HS2 leaks HMS connections
- [HIVE-8298](#) - Incorrect results for n-way join when join expressions are not in same order across joins
- [HIVE-8895](#) - bugs in mergejoin
- [HIVE-10771](#) - "separatorChar" has no effect in "CREATE TABLE AS SELECT" statement
- [HIVE-6679](#) - HiveServer2 should support configurable the server side socket timeout and keepalive for various transports types where applicable
- [HIVE-10732](#) - Hive JDBC driver does not close operation for metadata queries
- [HIVE-7027](#) - Hive job fails when referencing a view that explodes an array
- [IMPALA-1774](#) - Allow querying Parquet tables with complex-typed columns as long as those columns are not selected

- [IMPALA-1919](#) - Avoid calling ProcessBatch with out_batch->AtCapacity in right joins
- [IMPALA-2002](#) - Provide way to cache ext data source classes
- [IMPALA-1726](#) - Move JNI / Thrift utilities to separate header
- [HUE-2813](#) - [hive] Report when Hue server is down when trying to execute a query
- [HUE-2243](#) - [metastore] Listing tables can be very slow
- [OOZIE-1944](#) - Recursive variable resolution broken when same parameter name in config-default and action conf
- [PIG-4053](#) - TestMRCompiler succeeded with sun jdk 1.6 while failed with sun jdk 1.7
- [SENTRY-721](#) - HDFS Cascading permissions not applied to child file ACLs if a direct grant exists
- [SENTRY-699](#) - Memory leak when running Sentry w/ HiveServer2
- [SOLR-6146](#) - Leak in CloudSolrServer causing "Too many open files"
- [SOLR-7503](#) - Recovery after ZK session expiration happens in a single thread for all cores in a node

Issues Fixed in CDH 5.3.5

Potential job failures during YARN rolling upgrades to CDH 5.3.4

Problem: A MapReduce security fix introduced a compatibility issue that results in job failures during YARN rolling upgrades from CDH 5.3.3 to CDH 5.3.4.

Cloudera Bug: CDH-28680

Release affected: CDH 5.3.4

Release containing the fix: CDH 5.3.5

Workarounds: You can use any one of the following workarounds for this issue:

- Upgrade to CDH 5.3.5.
- Restart any jobs that might have failed during the upgrade.
- Explicitly set the version of MapReduce to be used so it is picked on a per-job basis.
 1. Update the YARN property, **MR Application Classpath** (`mapreduce.application.classpath`), either in Cloudera Manager or in the `mapred-site.xml` file. Remove all existing values and add a new entry: `<parcel-path>/lib/hadoop-mapreduce/*`, where `<parcel-path>` is the absolute path to the parcel installation. For example, the default installation path for the CDH 5.3.3 parcel would be: `/opt/cloudera/parcels/CDH-5.3.3-1.cdh5.3.3.p0.5/lib/hadoop-mapreduce/*`.
 2. Wait until jobs submitted with the above client configuration change have run to completion.
 3. Upgrade to CDH 5.3.4.
 4. Update the **MR Application Classpath** (`mapreduce.application.classpath`) property to point to the new CDH 5.3.4 parcel.

Do *not* delete the old parcel until after all jobs submitted prior to the upgrade have finished running.

Upstream Issues Fixed

The following upstream issue has been fixed in CDH 5.3.5:

- [YARN-3811](#) - NodeManager restarts could lead to application failures

Issues Fixed in CDH 5.3.4

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.3.4:

- [HDFS-7980](#) - Incremental BlockReport will dramatically slow down the startup of a namenode
- [HDFS-8380](#) - Always call addStoredBlock on blocks which have been shifted from one storage to another
- [HDFS-7645](#) - Rolling upgrade is restoring blocks from trash multiple times
- [HDFS-7869](#) - Inconsistency in the return information while performing rolling upgrade
- [HDFS-7340](#) - make rollingUpgrade start/finalize idempotent
- [HDFS-7312](#) - Update DistCp v1 to optionally not use tmp location (branch-1 only)

- [HDFS-7530](#) - Allow renaming of encryption zone roots
- [HDFS-7587](#) - Edit log corruption can happen if append fails with a quota violation
- [YARN-3485](#) - FairScheduler headroom calculation doesn't consider maxResources for Fifo and FairShare policies
- [YARN-3491](#) - PublicLocalizer#addResource is too slow.
- [YARN-3021](#) - YARN's delegation-token handling disallows certain trust setups to operate properly over DistCp
- [YARN-3241](#) - FairScheduler handles "invalid" queue names inconsistently
- [YARN-3022](#) - Expose Container resource information from NodeManager for monitoring
- [YARN-2984](#) - Metrics for container's actual memory usage
- [YARN-3465](#) - Use LinkedHashMap to preserve order of resource requests
- [MAPREDUCE-6339](#) - Job history file is not flushed correctly because isTimerActive flag is not set true when flushTimerTask is scheduled.
- [MAPREDUCE-5710](#) - Backport MAPREDUCE-1305 to branch-1
- [MAPREDUCE-6238](#) - MR2 can't run local jobs with -libjars command options which is a regression from MR1
- [MAPREDUCE-6076](#) - Zero map split input length combine with none zero map split input length may cause MR1 job hung sometimes.
- [HBASE-13374](#) - Small scanners (with particular configurations) do not return all rows
- [HBASE-13269](#) - Limit result array preallocation to avoid OOME with large scan caching values
- [HBASE-13422](#) - remove use of StandardCharsets in 0.98
- [HBASE-13335](#) - Update ClientSmallScanner and ClientSmallReversedScanner
- [HBASE-13262](#) - ResultScanner doesn't return all rows in Scan
- [HIVE-10646](#) - ColumnValue does not handle NULL_TYPE
- [HIVE-10453](#) - HS2 leaking open file descriptors when using UDFs
- [HIVE-9655](#) - Dynamic partition table insertion error
- [HIVE-10452](#) - Followup fix for HIVE-10202 to restrict it for script mode.
- [HIVE-10312](#) - SASL.QOP in JDBC URL is ignored for Delegation token Authentication
- [HIVE-10202](#) - Beeline outputs prompt+query on standard output when used in non-interactive mode
- [HIVE-10087](#) - Beeline's --silent option should suppress query from being echoed when running with -f option
- [HIVE-10085](#) - Lateral view on top of a view throws RuntimeException
- [HIVE-2828](#) - make timestamp accessible in the hbase KeyValue
- [HUE-2741](#) - [home] Hide the document move dialog
- [HUE-2732](#) - Hue isn't correctly doing add_column migrations with non-blank defaults
- [HUE-2513](#) - [fb] File list column sorting is broken
- [IMPALA-1519](#) - Fix wrapping of exprs via a TupleIsNullPredicate with analytics
- [IMPALA-1952](#) - Expand parsing of decimals to include scientific notation
- [IMPALA-1860](#) - INSERT/CTAS evaluates and applies constant predicates.
- [IMPALA-1900](#) - Assign predicates below analytic functions with a compatible partition by clause
- [IMPALA-1376](#) - Split up Planner into multiple classes.
- [IMPALA-1888](#) - FIRST_VALUE may produce incorrect results with preceding windows
- [IMPALA-1559](#) - FIRST_VALUE rewrite fn type might not match slot type
- [IMPALA-1808](#) - AnalyticEvalNode cannot handle partition/order by exprs with NaN
- [IMPALA-1562](#) - AnalyticEvalNode not properly handling nullable tuples
- [OOZIE-2063](#) - Cron syntax creates duplicate actions
- [OOZIE-2218](#) - META-INF directories in the war file have 777 permissions
- [OOZIE-1878](#) - Can't execute dryrun on the CLI
- [SENTRY-696](#) - Improve Metastoreplugin Cache Initialization time
- [SENTRY-703](#) - Calls to add_partition fail when passed a Partition object with a null location
- [SENTRY-408](#) - The URI permission should support more filesystem prefixes
- [SOLR-7478](#) - UpdateLog#close shutdown it's executor with interrupts before running close, preventing a clean close.
- [SOLR-7437](#) - Make HDFS transaction log replication factor configurable.

- [SOLR-7338](#) - A reloaded core will never register itself as active after a ZK session expiration
- [SOLR-7370](#) - FSHDFSUtils#recoverFileLease tries to recover the lease every one second after the first four second wait.
- [SPARK-6578](#) - Outbound channel in network library is not thread-safe, can lead to fetch failures
- [SQOOP-2343](#) - AsyncSqlRecordWriter stucks if any exception is thrown out in its close method
- [SQOOP-2286](#) - Ensure Sqoop generates valid avro column names
- [SQOOP-2283](#) - Support usage of --exec and --password-alias
- [SQOOP-2281](#) - Set overwrite on kite dataset
- [SQOOP-2282](#) - Add validation check for --hive-import and --append
- [SQOOP-2257](#) - Parquet target for imports with Hive overwrite option does not work
- [ZOOKEEPER-2146](#) - BinaryInputArchive readString should check length before allocating memory
- [ZOOKEEPER-2149](#) - Logging of client address when socket connection established

Published Known Issues Fixed

As a result of the above fixes, the following issues, previously published as [Known Issues in CDH 5](#) on page 1944, are also fixed.

Executing oozie job -config properties file -dryrun fails because of a code defect in argument parsing

Bug: [OOZIE-1878](#)

Severity: Low

Workaround: None.

Issues Fixed in CDH 5.3.3

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.3.3:

- [HADOOP-11722](#) - Some Instances of Services using ZKDelegationTokenSecretManager go down when old token cannot be deleted
- [HADOOP-11469](#) - KMS should skip default.key.acl and whitelist.key.acl when loading key acl
- [HADOOP-11710](#) - Make CryptoOutputStream behave like DFSOutputStream wrt synchronization
- [HADOOP-11674](#) - oneByteBuf in CryptoInputStream and CryptoOutputStream should be non static
- [HADOOP-11445](#) - Bzip2Codec: Data block is skipped when position of newly created stream is equal to start of split
- [HADOOP-11620](#) - Add support for load balancing across a group of KMS for HA
- [HDFS-6830](#) - BlockInfo.addStorage fails when DN changes the storage for a block replica
- [HDFS-7961](#) - Trigger full block report after hot swapping disk
- [HDFS-7960](#) - The full block report should prune zombie storages even if they're not empty
- [HDFS-7575](#) - Upgrade should generate a unique storage ID for each volume
- [HDFS-7596](#) - NameNode should prune dead storages from storageMap
- [HDFS-7579](#) - Improve log reporting during block report rpc failure
- [HDFS-7208](#) - NN doesn't schedule replication when a DN storage fails
- [HDFS-6899](#) - Allow changing MiniDFScluster volumes per DN and capacity per volume
- [HDFS-6878](#) - Change MiniDFScluster to support StorageType configuration for individual directories
- [HDFS-6678](#) - MiniDFScluster may still be partially running after initialization fails.
- [YARN-3351](#) - AppMaster tracking URL is broken in HA
- [YARN-3242](#) - Asynchrony in ZK-close can lead to ZKRMStateStore watcher receiving events for old client
- [YARN-2865](#) - Application recovery continuously fails with "Application with id already present. Cannot duplicate"
- [MAPREDUCE-6275](#) - Race condition in FileOutputCommitter v2 for user-specified task output subdirs
- [MAPREDUCE-4815](#) - Speed up FileOutputCommitter#commitJob for many output files
- [HBASE-13131](#) - ReplicationAdmin leaks connections if there's an error in the constructor
- [HIVE-10086](#) - Hive throws error when accessing Parquet file schema using field name match

- [HIVE-10098](#) - HS2 local task for map join fails in KMS encrypted cluster
- [HIVE-7426](#) - ClassCastException: ...IntWritable cannot be cast to ...Text involving ql.udf.generic.GenericUDFBasePad.evaluate
- [HIVE-7737](#) - Hive logs full exception for table not found
- [HIVE-9749](#) - ObjectStore schema verification logic is incorrect
- [HIVE-9788](#) - Make double quote optional in tsv/csv/dsv output
- [HIVE-9755](#) - Hive built-in "ngram" UDAF fails when a mapper has no matches.
- [HIVE-9770](#) - Beeline ignores --showHeader for non-tabular output formats i.e csv,tsv,dsv
- [HIVE-8688](#) - serialized plan OutputStream is not being closed
- [HIVE-9716](#) - Map job fails when table's LOCATION does not have scheme
- [HIVE-5857](#) - Reduce tasks do not work in uber mode in YARN
- [HIVE-8938](#) - Compiler should save the transform URI as input entity
- [HUE-2569](#) - [home] Delete project is broken
- [HUE-2529](#) - Increase the character limit of 'Name' Textfield in Useradmin Ldap Sync Groups
- [HUE-2506](#) - [search] Marker map does not display with HTML widget
- [HUE-1663](#) - [core] Option to either follow or not LDAP referrals for auth
- [HUE-2198](#) - [core] Reduce noise such as "handle_other(): Mutual authentication unavailable on 200 response"
- [SENTRY-683](#) - HDFS service client should ensure the kerberos ticket validity before new service connection
- [SENTRY-654](#) - Calls to append_partition fail when Sentry is enabled
- [SENTRY-664](#) - After Namenode is restarted, Path updates remain unsynched
- [SENTRY-665](#) - PathsUpdate.parsePath needs to handle special characters
- [SENTRY-652](#) - Sentry fails to parse spaces when HDFS ACL sync enabled
- [SOLR-7092](#) - Stop the HDFS lease recovery retries on HdfsTransactionLog on close and try to avoid lease recovery on closed files.
- [SOLR-7141](#) - RecoveryStrategy: Raise time that we wait for any updates from the leader before they saw the recovery state to have finished.
- [SOLR-7113](#) - Multiple calls to UpdateLog#init is not thread safe with respect to the HDFS FileSystem client object usage.
- [SOLR-7134](#) - Replication can still cause index corruption.
- [SQOOP-1764](#) - Numeric Overflow when getting extent map
- [IMPALA-1658](#) - Add compatibility flag for Hive-Parquet-Timestamps
- [IMPALA-1820](#) - Start with small pages for hash tables during repartitioning
- [IMPALA-1897](#) - Fixes for old hash join and agg
- [IMPALA-1894](#) - Fix old aggregation node hash table cleanup
- [IMPALA-1863](#) - Avoid deadlock across fragment instances
- [IMPALA-1915](#) - Fix query hang in BufferedBlockMgr:FindBlock()
- [IMPALA-1890](#) - Fixing a race between ~BufferedBlockMgr() and the WriteComplete() call
- [IMPALA-1738](#) - Use sprintf() instead of lexical_cast() in float-to-string casts
- [IMPALA-1865](#) - Fix partition spilling cleanup when new stream OOMs
- [IMPALA-1835](#) - Keep the fragment alive for TransmitData()
- [IMPALA-1805](#) - Impala's ACLs check do not consider all group ACLs, only checked first one.
- [IMPALA-1794](#) - Fix infinite loop opening or closing file with invalid metadata
- [IMPALA-1801](#) - external-data-source-executor leaking global jni refs
- [IMPALA-1712](#) - Unexpected remote bytes read counter was not being reset properly
- [IMPALA-1636](#) - Generalize index-based partition pruning to allow constant expressions

Published Known Issues Fixed

As a result of the above fixes, the following issues, previously published as [Known Issues in CDH 5](#) on page 1944, are also fixed.

After upgrade from a release earlier than CDH 5.2.0, storage IDs may no longer be unique

As of CDH 5.2, each storage volume on a DataNode should have its own unique `storageID`, but in clusters upgraded from CDH 4, or CDH 5 releases earlier than CDH 5.2.0, each volume on a given DataNode shares the same `storageID`, because the HDFS upgrade does not properly update the IDs to reflect the new naming scheme. This causes problems with load balancing. The problem affects only clusters upgraded from CDH 5.1.x and earlier to CDH 5.2 or later. Clusters that are new as of CDH 5.2.0 or later do not have the problem.

Bug: [HDFS-7575](#)

Cloudera Bug: CDH-24155

Severity: Medium

Workaround: Upgrade to a later or patched version of CDH.

Issues Fixed in CDH 5.3.2

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.3.2:

- [AVRO-1630](#) - Creating Builder from instance loses data
- [AVRO-1628](#) - Add Schema.createUnion(Schema... type)
- [AVRO-1539](#) - Add FileSystem-based FsInput Constructor
- [AVRO-1623](#) - GenericData#validate() of enum: IndexOutOfBoundsException
- [AVRO-1614](#) - Always getting a value...
- [AVRO-1592](#) - Java keyword as an enum constant in Avro schema file causes deserialization to fail.
- [AVRO-1619](#) - Generate better JavaDoc
- [AVRO-1622](#) - Add missing license headers
- [AVRO-1604](#) - ReflectData.AllowNull fails to generate schemas when @Nullable is present.
- [AVRO-1407](#) - NettyTransceiver can cause a infinite loop when slow to connect
- [AVRO-834](#) - Data File corruption recovery tool
- [AVRO-1596](#) - Cannot read past corrupted block in Avro data file
- [HADOOP-11350](#) - The size of header buffer of HttpServer is too small when HTTPS is enabled
- [HDFS-7707](#) - Edit log corruption due to delayed block removal again
- [HDFS-7718](#) - Store KeyProvider in ClientContext to avoid leaking key provider threads when using FileContext
- [HDFS-6425](#) - Large postponedMisreplicatedBlocks has impact on blockReport latency
- [HDFS-7560](#) - ACLs removed by removeDefaultAcl() will be back after NameNode restart/failover
- [HDFS-7513](#) - HDFS inotify: add defaultBlockSize to CreateEvent
- [HDFS-7158](#) - Reduce the memory usage of WebImageViewer
- [HDFS-7497](#) - Inconsistent report of decommissioning DataNodes between dfsadmin and NameNode webui
- [HDFS-6917](#) - Add an hdfs debug command to validate blocks, call recoverlease, etc.
- [HDFS-6779](#) - Add missing version subcommand for hdfs
- [YARN-2697](#) - RMAuthenticationHandler is no longer useful
- [YARN-2656](#) - RM web services authentication filter should add support for proxy user
- [YARN-3082](#) - Non thread safe access to systemCredentials in NodeHeartbeatResponse processing
- [YARN-3079](#) - Scheduler should also update maximumAllocation when updateNodeResource.
- [YARN-2992](#) - ZKRMStateStore crashes due to session expiry
- [YARN-2675](#) - containersKilled metrics is not updated when the container is killed during localization
- [YARN-2715](#) - Proxy user is problem for RPC interface if yarn.resourcemanager.webapp.proxyuser is not set
- [MAPREDUCE-6198](#) - NPE from JobTracker#resolveAndAddToTopology in MR1 cause initJob and heartbeat failure.
- [MAPREDUCE-6196](#) - Fix BigDecimal ArithmeticException in PiEstimator
- [HBASE-12540](#) - TestRegionServerMetrics#testMobMetrics test failure
- [HBASE-12533](#) - staging directories are not deleted after secure bulk load
- [HBASE-12077](#) - FilterLists create many ArrayList\$Itr objects per row.

- [HBASE-12386](#) - Replication gets stuck following a transient zookeeper error to remote peer cluster
- [HBASE-11979](#) - Compaction progress reporting is wrong
- [HBASE-12445](#) - hbase is removing all remaining cells immediately after the cell marked with marker = KeyValue.Type.DeleteColumn via PUT
- [HIVE-7647](#) - Beeline does not honor --headerInterval and --color when executing with "-e"
- [HIVE-7733](#) - Ambiguous column reference error on query
- [HIVE-9303](#) - Parquet files are written with incorrect definition levels
- [HIVE-8444](#) - update pom to junit 4.11
- [HIVE-9474](#) - truncate table changes permissions on the target
- [HIVE-9462](#) - HIVE-8577 - breaks type evolution
- [HIVE-9482](#) - Hive parquet timestamp compatibility
- [HIVE-6308](#) - COLUMNS_V2 Metastore table not populated for tables created without an explicit column list.
- [HIVE-9502](#) - Parquet cannot read Map types from files written with Hive 0.12 or earlier
- [HIVE-9445](#) - Revert HIVE-5700 - enforce single date format for partition column storage
- [HIVE-9393](#) - reduce noisy log level of ColumnarSerDe.java:116 from INFO to DEBUG
- [HIVE-7800](#) - Parquet Column Index Access Schema Size Checking
- [HIVE-9330](#) - DummyTxnManager will throw NPE if WriteEntity writeType has not been set
- [HIVE-9265](#) - Hive with encryption throws NPE to fs path without schema
- [HIVE-9199](#) - Excessive exclusive lock used in some DDLs with DummyTxnManager
- [HIVE-6978](#) - beeline always exits with 0 status, should exit with non-zero status on error
- [HUE-2556](#) - [core] Cannot update project tags of a document
- [HUE-2528](#) - Partitions limit gets capped to 1000 despite configuration
- [HUE-2548](#) - [metastore] Create table then load data does redirect to the table page
- [HUE-2525](#) - [core] Fix manual install of samples
- [HUE-2501](#) - [metastore] Creating a table with header files bigger than 64MB truncates it
- [HUE-2484](#) - [beeswax] Configure support for Hive Server2 LDAP authentication
- [HUE-2532](#) - [search] Fix share URL on Internet Explorer
- [HUE-2531](#) - [impala] Autogrow missing result list
- [HUE-2524](#) - [impala] Sort numerically recent queries tab
- [HUE-2495](#) - [oozie] Improve dashboards sorting mechanism
- [HUE-2511](#) - [impala] Infinite scroll keeps fetching results even if finished
- [HUE-2102](#) - [oozie] Workflow with credentials can't be used with Coordinator
- [HUE-2152](#) - [pig] Credentials support in editor
- [OOZIE-2131](#) - Add flag to sqoop action to skip hbase delegation token generation
- [OOZIE-2047](#) - Oozie does not support Hive tables that use datatypes introduced since Hive 0.8
- [OOZIE-2102](#) - Streaming actions are broken cause of incorrect method signature
- [PARQUET-173](#) - StatisticsFilter doesn't handle And properly
- [PARQUET-157](#) - Divide by zero in logging code
- [PARQUET-142](#) - parquet-tools doesn't filter _SUCCESS file
- [PARQUET-124](#) - parquet.hadoop.ParquetOutputCommitter.commitJob() throws parquet.io.ParquetEncodingException
- [PARQUET-136](#) - NPE thrown in StatisticsFilter when all values in a string/binary column trunk are null
- [PARQUET-168](#) - Wrong command line option description in parquet-tools
- [PARQUET-145](#) - InternalParquetRecordReader.close() should not throw an exception if initialization has failed
- [PARQUET-140](#) - Allow clients to control the GenericData object that is used to read Avro records
- [SOLR-7033](#) - [RecoveryStrategy should not publish any state when closed / cancelled.
- [SOLR-5961](#) - Solr gets crazy on /overseer/queue state change
- [SOLR-6640](#) - Replication can cause index corruption
- [SOLR-5875](#) - QueryComponent.mergelds() unmarshals all docs' sort field values once per doc instead of once per shard

- [SOLR-6919](#) - Log REST info before executing
- [SOLR-6969](#) - When opening an HDFSTransactionLog for append we must first attempt to recover it's lease to prevent data loss.
- [SOLR-5515](#) - NPE when getting stats on date field with empty result on solrcloud
- [SPARK-3778](#) - newAPIHadoopRDD doesn't properly pass credentials for secure hdfs on yarn
- [SPARK-4835](#) - Streaming saveAs*HadoopFiles() methods may throw FileAlreadyExistsException during checkpoint recovery
- [SQOOP-2057](#) - Skip delegation token generation flag during hbase import
- [SQOOP-1779](#) - Add support for --hive-database when importing Parquet files into Hive
- [IMPALA-1622](#) - Fix overflow in StringParser::StringToFloatInternal()
- [IMPALA-1614](#) - Compute stats fails if table name starts with number
- [IMPALA-1623](#) - unix_timestamp() does not return correct time
- [IMPALA-1535](#) - Partition pruning with NULL
- [IMPALA-1606](#) - Impala does not always give short name to Llama
- [IMPALA-1120](#) - Fetch column statistics using Hive 0.13 bulk API

In addition, CDH 5.3.2 reverts [YARN-2713](#), which has caused problems since its inclusion in CDH 5.3.0.

Published Known Issues Fixed

As a result of the above fixes, the following issues, previously published as [Known Issues in CDH 5](#) on page 1944, are also fixed.

Hive does not support Parquet schema evolution

Adding a new column to a Parquet table causes queries on that table to fail with a `column not found` error.

Bug: [HIVE-7800](#)

Cloudera Bug: CDH-24123

Severity: Medium

Workaround: Use Impala instead; Impala handles Parquet schema evolution correctly.

Issues Fixed in CDH 5.3.1

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.3.1:

- [YARN-2975](#) - FSLeafQueue app lists are accessed without required locks
- [YARN-2010](#) - Handle app-recovery failures gracefully
- [YARN-3027](#) - Scheduler should use totalAvailable resource from node instead of availableResource for maxAllocation
- [HIVE-9445](#) - Revert HIVE-5700 - enforce single date format for partition column storage
- [IMPALA-1668](#) - TSaslServerTransport::Factory::getTransport() leaks transport map entries
- [IMPALA-1674](#) - IMPALA-1556 causes memory leak with secure connections

Published Known Issues Fixed

As a result of the above fixes, the following issues, previously published as [Known Issues in CDH 5](#) on page 1944, are also fixed.

Upgrading a PostgreSQL Hive Metastore from Hive 0.12 to Hive 0.13 may result in a corrupt metastore

[HIVE-5700](#) introduced a serious bug into the Hive Metastore upgrade scripts. This bug affects users who have a PostgreSQL Hive Metastore and have *at least one table* which is partitioned by date and the value is stored as a date type (not string).

Bug: [HIVE-5700](#)

Cloudera Bug: CDH-24531

Severity: High

Workaround: None. Do not upgrade your PostgreSQL metastore to version 0.13 if you satisfy the condition stated above.

Issues Fixed in CDH 5.3.0

The following topics describe known issues fixed in CDH 5.3.0.

Apache Hadoop

HDFS

Kerberos re-login attempts fail when using JDK 1.7.0_80

On clusters using JDK 1.7.0_80, long running HDFS clients are unable to re-authenticate using Kerberos once their ticket expires. Due to this authentication failure, any jobs triggered from these clients will fail.

Releases Affected: CDH 5.1, 5.2

Bug: [HADOOP-10786](#)

Cloudera Bug: CDH-27714

Workaround: Upgrade to CDH 5.3.2 (or higher).

NameNode - KMS communication fails after long periods of inactivity

Encrypted files and encryption zones cannot be created if a long period of time (by default, 20 hours) has passed since the last time the KMS and NameNode communicated.

Bug: [HADOOP-11187](#)

Cloudera Bug: CDH-22217

Workaround: There are two possible workarounds to this issue:

- You can increase the KMS authentication token validity period to a very high number. Since the default value is 10 hours, this bug will only be encountered after 20 hours of no communication between the NameNode and the KMS. Add the following property to the `kms-site.xml` Safety Valve:

```
<property>
<name>hadoop.kms.authentication.token.validity</name>
<value>SOME VERY HIGH NUMBER</value>
</property>
```

- You can switch the KMS signature secret provider to the string secret provider by adding the following property to the `kms-site.xml` Safety Valve:

```
<property>
<name>hadoop.kms.authentication.signature.secret</name>
<value>SOME VERY SECRET STRING</value>
</property>
```

DataNodes may become unresponsive to block creation requests

In releases earlier than CDH 5.2.3, DataNodes may become unresponsive to block creation requests from clients when the directory scanner is running.

Bug: [HDFS-7489](#)

Cloudera Bug: CDH-23732

Workaround: Upgrade to CDH 5.2.3 or later.

Apache Hive

UDF translate() does not accept arguments of type CHAR or VARCHAR

Bug: [HIVE-6622](#)

Workaround: Cast the argument to type String.

Hive's Timestamp type cannot be stored in Parquet

Tables containing timestamp columns cannot use Parquet as the storage engine.

Bug: [HIVE-6394](#)

Workaround: Use a different file format.

Apache Spark

Spark sort-based shuffle is affected by a kernel bug

Spark sort-based shuffle is affected by a kernel bug

(<http://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/commit/?id=2cb4b05e7647891b46b91c07c9a60304803d1688>).

The kernel bug was fixed in RHEL/CentOS 6.2.



Note: Previously CDH defaulted to hash-based shuffle. It now defaults to sort-based shuffle.

Bug: [SPARK-3948](#)

Cloudera Bug: CDH-23577

Issues Fixed in CDH 5.2.x

The following topics describe known issues fixed in CDH 5.2.x, from newest to oldest release.

CDH 5.2.6

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.2.6:

- [CRUNCH-516](#) - Scrunch needs some additional null checks
- [CRUNCH-508](#) - Improve performance of Scala Enumeration counters in Scrunch
- [CRUNCH-514](#) - AvroDerivedDeepCopier should initialize delegate MapFns
- [CRUNCH-530](#) - Fix object reuse bug in GenericRecordToTuple
- [HADOOP-12103](#) - Small refactoring of DelegationTokenAuthenticationFilter to allow code sharing
- [HADOOP-10839](#) - Add unregisterSource() to MetricsSystem API
- [HDFS-8337](#) - Accessing httpfs via webhdfs doesn't work from a jar with kerberos
- [HDFS-7546](#) - Document, and set an accepting default for dfs.namenode.kerberos.principal.pattern
- [HDFS-6997](#) - Archival Storage: add more tests for data migration and replicaion
- [HDFS-7980](#) - Incremental BlockReport will dramatically slow down the startup of a namenode
- [HDFS-8380](#) - Always call addStoredBlock on blocks which have been shifted from one storage to another
- [HDFS-7312](#) - Update DistCp v1 to optionally not use tmp location (branch-1 only)
- [YARN-3485](#) - FairScheduler headroom calculation doesn't consider maxResources for Fifo and FairShare policies
- [YARN-3241](#) - FairScheduler handles "invalid" queue names inconsistently
- [YARN-2669](#) - FairScheduler: queue names shouldn't allow periods
- [YARN-3022](#) - Expose Container resource information from NodeManager for monitoring
- [YARN-2984](#) - Metrics for container's actual memory usage
- [YARN-3465](#) - Use LinkedHashMap to preserve order of resource requests
- [MAPREDUCE-6387](#) - Serialize the recently added Task#encryptedSpillKey field at the end
- [MAPREDUCE-6339](#) - Job history file is not flushed correctly because isTimerActive flag is not set true when flushTimerTask is scheduled.
- [MAPREDUCE-5710](#) - Backport MAPREDUCE-1305 to branch-1
- [MAPREDUCE-6238](#) - MR2 can't run local jobs with -libjars command options which is a regression from MR1
- [HBASE-13826](#) - Unable to create table when group acls are appropriately set.
- [HBASE-13241](#) - Add tests for group level grants
- [HBASE-13239](#) - HBase grant at specific column level does not work for Groups

- [HBASE-13768](#) - ZooKeeper znodes are bootstrapped with insecure ACLs in a secure configuration
- [HBASE-13789](#) - ForeignException should not be sent to the client
- [HBASE-13779](#) - Calling table.exists() before table.get() end up with an empty Result
- [HBASE-13780](#) - Default to 700 for HDFS root dir permissions for secure deployments
- [HBASE-13768](#) - ZooKeeper znodes are bootstrapped with insecure ACLs in a secure configuration
- [HBASE-13767](#) - Allow ZKAcIReset to set and not just clear ZK ACLs
- [HBASE-13086](#) - Show ZK root node on Master WebUI
- [HBASE-13342](#) - Fix incorrect interface annotations
- [HBASE-13162](#) - Add capability for cleaning hbase acls to hbase cleanup script.
- [HBASE-12641](#) - Grant all permissions of hbase zookeeper node to hbase superuser in a secure cluster
- [HBASE-13374](#) - Small scanners (with particular configurations) do not return all rows
- [HBASE-13269](#) - Limit result array preallocation to avoid OOME with large scan caching values
- [HBASE-13422](#) - remove use of StandardCharsets in 0.98
- [HBASE-13335](#) - Update ClientSmallScanner and ClientSmallReversedScanner
- [HBASE-13262](#) - ResultScanner doesn't return all rows in Scan
- [HIVE-10841](#) - [WHERE col is not null] does not work sometimes for queries with many JOIN statements
- [HIVE-9620](#) - Cannot retrieve column statistics using HMS API if column name contains uppercase characters
- [HIVE-8863](#) - Cannot drop table with uppercase name after "compute statistics for columns"
- [HIVE-6679](#) - HiveServer2 should support configurable the server side socket timeout and keepalive for various transports types where applicable
- [OOZIE-1944](#) - Recursive variable resolution broken when same parameter name in config-default and action conf
- [OOZIE-2218](#) - META-INF directories in the war file have 777 permissions
- [SENTRY-540](#) - Fix Sentry test validating special chars in username due to HIVE-8916
- [SENTRY-227](#) - Fix for "Unsupported entity type DUMMYPARTITION"
- [SOLR-7478](#) - UpdateLog#close shutdown it's executor with interrupts before running close, preventing a clean close.
- [SOLR-7338](#) - A reloaded core will never register itself as active after a ZK session expiration
- [SOLR-7370](#) - FSHDFSUtils#recoverFileLease tries to recover the lease every one second after the first four second wait.
- [ZOOKEEPER-2146](#) - BinaryInputArchive readString should check length before allocating memory
- [ZOOKEEPER-2149](#) - Logging of client address when socket connection established
- [IMPALA-1726](#): Move JNI / Thrift utilities to separate header
- [IMPALA-2002](#): Provide way to cache ext data source classes

CDH 5.2.5

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.2.5:

- [HADOOP-11350](#) - The size of header buffer of HttpServer is too small when HTTPS is enabled
- [HADOOP-11710](#) - Make CryptoOutputStream behave like DFSOutputStream wrt synchronization
- [HADOOP-11674](#) - oneByteBuf in CryptoInputStream and CryptoOutputStream should be non static
- [HDFS-6830](#) - BlockInfo.addStorage fails when DN changes the storage for a block replica
- [HDFS-7960](#) - The full block report should prune zombie storages even if they're not empty
- [HDFS-6425](#) - Large postponedMisreplicatedBlocks has impact on blockReport latency
- [HDFS-7575](#) - Upgrade should generate a unique storage ID for each volume
- [HDFS-7596](#) - NameNode should prune dead storages from storageMap
- [HDFS-7579](#) - Improve log reporting during block report rpc failure
- [HDFS-7208](#) - NN doesn't schedule replication when a DN storage fails
- [HDFS-6899](#) - Allow changing MiniDFScluster volumes per DN and capacity per volume
- [HDFS-6878](#) - Change MiniDFScluster to support StorageType configuration for individual directories
- [HDFS-6678](#) - MiniDFScluster may still be partially running after initialization fails.

- [HDFS-7575](#) - Upgrade should generate a unique storage ID for each volume
- [HDFS-7960](#) - The full block report should prune zombie storages even if they're not empty
- [YARN-570](#) - Time strings are formatted in different timezone
- [YARN-2251](#) - Avoid negative elapsed time in JHS/MRAM web UI and services
- [YARN-3242](#) - Asynchrony in ZK-close can lead to ZKRMStateStore watcher receiving events for old client
- [MAPREDUCE-5957](#) - AM throws ClassNotFoundException with job classloader enabled if custom output format/committer is used
- [MAPREDUCE-6076](#) - Zero map split input length combine with none zero map split input length may cause MR1 job hung sometimes.
- [MAPREDUCE-6275](#) - Race condition in FileOutputCommitter v2 for user-specified task output subdirs
- [MAPREDUCE-4815](#) - Speed up FileOutputCommitter#commitJob for many output files
- [HIVE-2828](#) - make timestamp accessible in the hbase KeyValue
- [HIVE-2828](#) - make timestamp accessible in the hbase KeyValue
- [HIVE-7433](#) - ColumnMappers.ColumnMapping should expose public accessors for its fields
- [HIVE-6148](#) - Support arbitrary structs stored in HBase
- [HIVE-6147](#) - Support avro data stored in HBase columns
- [HIVE-6584](#) - Add HiveHBaseTableSnapshotInputFormat
- [HIVE-6411](#) - Support more generic way of using composite key for HBaseHandler
- [HIVE-6677](#) - HBaseSerDe needs to be refactored
- [HIVE-9934](#) - Vulnerability in LdapAuthenticationProviderImpl enables HiveServer2 client to degrade the authentication mechanism to "none", allowing authentication without password
- [HIVE-7737](#) - Hive logs full exception for table not found
- [HIVE-9716](#) - Map job fails when table's LOCATION does not have scheme
- [HIVE-8688](#) - serialized plan OutputStream is not being closed
- [HIVE-5857](#) - Reduce tasks do not work in uber mode in YARN
- [HUE-2446](#) - Migrating from CDH 4.7 to CDH 5.0.1+/Hue 3.5+ will fail
- [HUE-2371](#) - [sentry] Sentry URI should be created only with a ALL permission
- [HUE-1663](#) - [core] Option to either follow or not LDAP referrals for auth
- [SENTRY-654](#) - Calls to append_partition fail when Sentry is enabled
- [SOLR-7092](#) - Stop the HDFS lease recovery retries on HdfsTransactionLog on close and try to avoid lease recovery on closed files.
- [SOLR-7134](#) - Replication can still cause index corruption.
- [SOLR-7113](#) - Multiple calls to UpdateLog#init is not thread safe with respect to the HDFS FileSystem client object usage.
- [SOLR-7141](#) - RecoveryStrategy: Raise time that we wait for any updates from the leader before they saw the recovery state to have finished.
- [SQOOP-1764](#) - Numeric Overflow when getting extent map
- [IMPALA-1658](#): Add compatibility flag for Hive-Parquet-Timestamps
- [IMPALA-1794](#): Fix infinite loop opening/closing file w/ invalid metadata
- [IMPALA-1801](#): external-data-source-executor leaking global jni refs

CDH 5.2.4

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.2.4:

- [HDFS-7707](#) - Edit log corruption due to delayed block removal again
- [YARN-2846](#) - Incorrect persist exit code for running containers in reacquireContainer() that interrupted by NodeManager restart.
- [HIVE-7733](#) - Ambiguous column reference error on query
- [HIVE-8444](#) - update pom to junit 4.11
- [HIVE-9474](#) - truncate table changes permissions on the target

- [HIVE-6308](#) - COLUMNS_V2 Metastore table not populated for tables created without an explicit column list.
- [HIVE-9445](#) - Revert HIVE-5700 - enforce single date format for partition column storage
- [HIVE-7800](#) - Parquet Column Index Access Schema Size Checking Checking
- [HIVE-9393](#) - reduce noisy log level of ColumnarSerDe.java:116 from INFO to DEBUG
- [HUE-2501](#) - [metastore] Creating a table with header files bigger than 64MB truncates it
- [SOLR-7033](#) - [RecoveryStrategy should not publish any state when closed / cancelled.
- [SOLR-5961](#) - Solr gets crazy on /overseer/queue state change
- [SOLR-6640](#) - Replication can cause index corruption
- [SOLR-6920](#) - During replication use checksums to verify if files are the same
- [SOLR-5875](#) - QueryComponent.mergelds() unmarshals all docs' sort field values once per doc instead of once per shard
- [SOLR-6919](#) - Log REST info before executing
- [SOLR-6969](#) - When opening an HDFSTransactionLog for append we must first attempt to recover its lease to prevent data loss
- [IMPALA-1471](#): Bug in spilling of PHJ that was affecting left anti and outer joins.
- [IMPALA-1451](#): Empty Row in HBase triggers NPE in Planner
- [IMPALA-1535](#): Partition pruning with NULL
- [IMPALA-1483](#): Substitute TuplelsNullPredicates to refer to physical analytic output.
- [IMPALA-1674](#): Fix serious memory leak in TSaslTransport
- [IMPALA-1668](#): Fix leak of transport objects in TSaslServerTransport::Factory
- [IMPALA-1565](#): Python sasl client transport perf issue
- [IMPALA-1556](#): Kerberos fetches 3x slower
- [IMPALA-1120](#): Fetch column statistics using Hive 0.13 bulk API

Published Known Issues Fixed

As a result of the above fixes, the following issues, previously published as [Known Issues in CDH 5](#) on page 1944, are also fixed:

Hive does not support Parquet schema evolution

Adding a new column to a Parquet table causes queries on that table to fail with a `column not found` error.

Bug: [HIVE-7800](#)

Cloudera Bug: CDH-24123

Workaround: Use Impala instead; Impala handles Parquet schema evolution correctly.

CDH 5.2.3

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.2.3:

- [AVRO-1623](#) - GenericData#validate() of enum: IndexOutOfBoundsException
- [AVRO-1622](#) - Add missing license headers
- [AVRO-1604](#) - ReflectData.AllowNull fails to generate schemas when @Nullable is present.
- [AVRO-1407](#) - NettyTransceiver can cause a infinite loop when slow to connect
- [AVRO-834](#) - Data File corruption recovery tool
- [AVRO-1596](#) - Cannot read past corrupted block in Avro data file
- [CRUNCH-480](#) - AvroParquetFileSource does not properly configure user-supplied read schema
- [CRUNCH-479](#) - Writing to target with WriteMode.APPEND merges values into PCollection
- [CRUNCH-477](#) - Fix HFileTargetIT failures on hadoop1 under Java 1.7/1.8
- [CRUNCH-473](#) - Use specific class type for case class serialization
- [CRUNCH-473](#) - Use specific class type for case class serialization
- [CRUNCH-472](#) - Add Scrunch serialization support for Java Enums
- [HADOOP-11068](#) - Match hadoop.auth cookie format to jetty output

- [HADOOP-11343](#) - Overflow is not properly handled in calculating final iv for AES CTR
- [HADOOP-11301](#) - [optionally] update jmx cache to drop old metrics
- [HADOOP-11085](#) - Excessive logging by org.apache.hadoop.util.Progress when value is NaN
- [HADOOP-11247](#) - Fix a couple javac warnings in NFS
- [HADOOP-11195](#) - Move Id-Name mapping in NFS to the hadoop-common area for better maintenance
- [HADOOP-11130](#) - NFS updateMaps OS check is reversed
- [HADOOP-10990](#) - Add missed NFSv3 request and response classes
- [HADOOP-11323](#) - WritableComparator#compare keeps reference to byte array
- [HDFS-7560](#) - ACLs removed by removeDefaultAcl() will be back after NameNode restart/failover
- [HDFS-7367](#) - HDFS short-circuit read cannot negotiate shared memory slot and file descriptors when SASL is enabled on DataTransferProtocol.
- [HDFS-7489](#) - Incorrect locking in FsVolumeList#checkDirs can hang datanodes
- [HDFS-7158](#) - Reduce the memory usage of WebImageViewer
- [HDFS-7497](#) - Inconsistent report of decommissioning DataNodes between dfsadmin and NameNode webui
- [HDFS-7146](#) - NFS ID/Group lookup requires SSSD enumeration on the server
- [HDFS-7387](#) - NFS may only do partial commit due to a race between COMMIT and write
- [HDFS-7356](#) - Use DirectoryListing.hasMore() directly in nfs
- [HDFS-7180](#) - NFSv3 gateway frequently gets stuck due to GC
- [HDFS-7259](#) - Unresponsive NFS mount point due to deferred COMMIT response
- [HDFS-6894](#) - Add XDR parser method for each NFS response
- [HDFS-6850](#) - Move NFS out of order write unit tests into TestWrites class
- [HDFS-7385](#) - ThreadLocal used in FSEditLog class causes FSImage permission mess up
- [HDFS-7409](#) - Allow dead nodes to finish decommissioning if all files are fully replicated
- [HDFS-7373](#) - Clean up temporary files after fsimage transfer failures
- [HDFS-7225](#) - Remove stale block invalidation work when DN re-registers with different UUID
- [YARN-2721](#) - Race condition: ZKRMStateStore retry logic may throw NodeExist exception
- [YARN-2975](#) - FSLeafQueue app lists are accessed without required locks
- [YARN-2992](#) - ZKRMStateStore crashes due to session expiry
- [YARN-2910](#) - FSLeafQueue can throw ConcurrentModificationException
- [YARN-2816](#) - NM fail to start with NPE during container recovery
- [MAPREDUCE-6198](#) - NPE from JobTracker#resolveAndAddToTopology in MR1 cause initJob and heartbeat failure.
- [MAPREDUCE-6169](#) - MergeQueue should release reference to the current item from key and value at the end of the iteration to save memory.
- [HBASE-11794](#) - StripeStoreFlusher causes NullPointerException
- [HBASE-12077](#) - FilterLists create many ArrayList\$Itr objects per row.
- [HBASE-12386](#) - Replication gets stuck following a transient zookeeper error to remote peer cluster
- [HBASE-11979](#) - Compaction progress reporting is wrong
- [HBASE-12529](#) - Use ThreadLocalRandom for RandomQueueBalancer
- [HBASE-12445](#) - hbase is removing all remaining cells immediately after the cell marked with marker = KeyValue.Type.DeleteColumn via PUT
- [HBASE-12460](#) - Moving Chore to hbase-common module.
- [HBASE-12366](#) - Add login code to HBase Canary tool.
- [HBASE-12447](#) - Add support for setTimeRange for RowCounter and CellCounter
- [HIVE-9330](#) - DummyTxnManager will throw NPE if WriteEntity writeType has not been set
- [HIVE-9199](#) - Excessive exclusive lock used in some DDLs with DummyTxnManager
- [HIVE-6835](#) - Reading of partitioned Avro data fails if partition schema does not match table schema
- [HIVE-6978](#) - beeline always exits with 0 status, should exit with non-zero status on error
- [HIVE-8891](#) - Another possible cause to NucleusObjectNotFoundException from drops/rollback
- [HIVE-8874](#) - Error Accessing HBase from Hive via Oozie on Kerberos 5.0.1 cluster
- [HIVE-8916](#) - Handle user@domain username under LDAP authentication

- [HIVE-8889](#) - JDBC Driver ResultSet.getXXXXXX(String columnName) methods Broken
- [HIVE-9445](#) - Revert HIVE-5700 - enforce single date format for partition column storage
- [HIVE-5454](#) - HCatalog runs a partition listing with an empty filter
- [HIVE-8784](#) - Querying partition does not work with JDO enabled against PostgreSQL
- [HUE-2484](#) - [beeswax] Configure support for Hive Server2 LDAP authentication
- [HUE-2102](#) - [oozie] Workflow with credentials can't be used with Coordinator
- [HUE-2152](#) - [pig] Credentials support in editor
- [HUE-2472](#) - [impala] Stabilize result retrieval
- [HUE-2406](#) - [search] New dashboard page has a margin problem
- [HUE-2373](#) - [search] Heatmap can break
- [HUE-2395](#) - [search] Broken widget in Solr Apache logs example
- [HUE-2414](#) - [search] Timeline chart breaks when there's no extraSeries defined
- [HUE-2342](#) - [impala] TLS/SSL encryption
- [HUE-2426](#) - [pig] Dashboard gives a 500 error
- [HUE-2430](#) - [pig] Progress bars of running scripts not updated on Dashboard
- [HUE-2411](#) - [useradmin] Lazy load user and group list in permission sharing popup
- [HUE-2398](#) - [fb] Drag and Drop hover message should not appear when elements originating in DOM are dragged
- [HUE-2401](#) - [search] Visually report selected and excluded values for ranges too
- [HUE-2389](#) - [impala] Expand results table after the results are added to datatables
- [HUE-2360](#) - [sentry] Sometimes Groups are not loaded we see the input box instead
- [IMPALA-1453](#) - Fix many bugs with HS2 FETCH_FIRST
- [IMPALA-1623](#) - unix_timestamp() does not return correct time
- [IMPALA-1606](#) - Impala does not always give short name to Llama
- [IMPALA-1475](#) - accept unmangled native UDF symbols
- [OOZIE-2102](#) - Streaming actions are broken cause of incorrect method signature
- [PARQUET-145](#) - InternalParquetRecordReader.close() should not throw an exception if initialization has failed
- [PARQUET-140](#) - Allow clients to control the GenericData object that is used to read Avro records
- [PIG-4330](#) - Regression test for PIG-3584 - AvroStorage does not correctly translate arrays of strings
- [PIG-3584](#) - AvroStorage does not correctly translate arrays of strings
- [SOLR-5515](#) - NPE when getting stats on date field with empty result on solrcloud

Published Known Issues Fixed

As a result of the above fixes, the following issues, previously published as [Known Issues in CDH 5](#) on page 1944, are also fixed.

Upgrading a PostgreSQL Hive Metastore from Hive 0.12 to Hive 0.13 may result in a corrupt metastore

[HIVE-5700](#) introduced a serious bug into the Hive Metastore upgrade scripts. This bug affects users who have a PostgreSQL Hive Metastore and have *at least one table* which is partitioned by date and the value is stored as a date type (not string).

Bug: [HIVE-5700](#)

Cloudera Bug: CDH-24531

Workaround: None. Do not upgrade your PostgreSQL metastore to version 0.13 if you satisfy the condition stated above.

DataNodes may become unresponsive to block creation requests

DataNodes may become unresponsive to block creation requests from clients when the directory scanner is running.

Bug: [HDFS-7489](#)

Cloudera Bug: CDH-23732

Workaround: Disable the directory scanner by setting `dfs.datanode.directoryscan.interval` to `-1`.

CDH 5.2.2

There is no CDH 5.2.2 release.

CDH 5.2.1

Apache Hadoop

Files inside encryption zones cannot be read in Hue

Hue uses either WebHDFS or HttpFS to access files. Both are proxy user clients of KMS and the KMS client library does not currently handle proxy users correctly.

Bug: [HADOOP-11176](#)

Cloudera Bug: CDH-22276

Workaround: None

Both ResourceManagers can end up in Standby mode

After a restart, if an application fails to recover, both Resource Managers could end up in Standby mode.

Bug: [YARN-2588](#), [YARN-2010](#)

Cloudera Bug: CDH-22337

Workarounds:

- Stop the RM. Format the state store using `yarn resourcemanager -format-state-store`. Applications that were running before the ResourceManager went down will not be recovered.
- You can limit the number of completed applications the RM state-store stores (`yarn.resourcemanager.state-store.max-completed-applications`) to reduce the chances of running into this problem.

Apache Oozie

Using cron-like syntax for Coordinator frequencies can result in duplicate actions

Every `throttle` number of actions will be a duplicate. For example, if the throttle is set to 5, every fifth action will be a duplicate.

Bug: [OOZIE-2063](#)

Cloudera Bug: CDH-23137

Workaround: If possible, use the older syntax to specify an equivalent frequency.

Upstream Fixes

CDH 5.2.1 also fixes the following issues:

- [HADOOP-11243](#) - SSLFactory shouldn't allow SSLv3
- [HADOOP-11217](#) - Disable SSLv3 in KMS
- [HADOOP-11156](#) - DelegateToFileSystem should implement `getFsStatus(final Path f)`.
- [HADOOP-11176](#) - KMSSClientProvider authentication fails when both `currentUgi` and `loginUgi` are a proxied user
- [HDFS-7235](#) - `DataNode#transferBlock` should report blocks that don't exist using `reportBadBlock`
- [HDFS-7274](#) - Disable SSLv3 in HttpFS
- [HDFS-7391](#) - Reenable `SSLv2Hello` in HttpFS
- [HDFS-6781](#) - Separate HDFS commands from `CommandsManual.apptm`
- [HDFS-6831](#) - Inconsistency between `'hdfs dfsadmin'` and `'hdfs dfsadmin -help'`
- [HDFS-7278](#) - Add a command that allows sysadmins to manually trigger full block reports from a DN
- [YARN-2010](#) - Handle app-recovery failures gracefully
- [YARN-2588](#) - Standby RM does not `transitionToActive` if previous `transitionToActive` is failed with ZK exception.
- [YARN-2566](#) - `DefaultContainerExecutor` should pick a working directory randomly
- [YARN-2641](#) - Decommission nodes on `-refreshNodes` instead of next NM-RM heartbeat
- [MAPREDUCE-6147](#) - Support `mapreduce.input.fileinputformat.split.maxsize`

- [HBASE-12376](#) - HBaseAdmin leaks ZK connections if failure starting watchers (ConnectionLossException)
- [HBASE-12201](#) - Close the writers in the MOB sweep tool
- [HBASE-12220](#) - Add hedgedReads and hedgedReadWins metrics
- [HIVE-8693](#) - Separate out fair scheduler dependency from hadoop 0.23 shim
- [HIVE-8634](#) - HiveServer2 fair scheduler queue mapping does not handle the secondary groups rules correctly
- [HIVE-8675](#) - Increase thrift server protocol test coverage
- [HIVE-8827](#) - Remove SSLv2Hello from list of disabled protocols protocols
- [HIVE-8615](#) - beeline csv,tsv outputformat needs backward compatibility mode
- [HIVE-8627](#) - Compute stats on a table from impala caused the table to be corrupted
- [HIVE-7764](#) - Support all JDBC-HiveServer2 authentication modes on a secure cluster cluster
- [HIVE-8182](#) - beeline fails when executing multiple-line queries with trailing spaces
- [HUE-2438](#) - [core] Disable SSLv3 for Poodle vulnerability
- [IMPALA-1361](#): FE Exceptions with BETWEEN predicates
- [IMPALA-1397](#): free local expr allocations in scanner threads
- [IMPALA-1400](#): Window function insert issue (LAG() + OVER)
- [IMPALA-1401](#): raise MAX_PAGE_HEADER_SIZE and use scanner context to stitch together header buffer
- [IMPALA-1410](#): accept "single character" character classes in regex functions
- [IMPALA-1411](#): Create table as select produces incorrect results
- [IMPALA-1416](#) - Queries fail with metastore exception after upgrade and compute stats
- [OOZIE-2034](#) - Disable SSLv3 (POODLEbleed vulnerability)
- [OOZIE-2063](#) - Cron syntax creates duplicate actions
- [PARQUET-107](#) - Add option to disable summary metadata aggregation after MR jobs
- [SPARK-3788](#) - Yarn dist cache code is not friendly to HDFS HA, Federation
- [SPARK-3661](#) - spark.*.memory is ignored in cluster mode
- [SPARK-3979](#) - Yarn backend's default file replication should match HDFS' default one
- [SPARK-1720](#) - use LD_LIBRARY_PATH instead of -Djava.library.path

CDH 5.2.0

Apache HBase

hbase.zookeeper.useMulti set to false by default

The default value of `hbase.zookeeper.useMulti` was changed from `true` to `false` in CDH 5. In CDH 5.2, the default is changed back to `true`. This affects environments with HBase replication enabled and large replication queues.

Cloudera Bug: CDH-22716

Workaround: Enable `hbase.zookeeper.useMulti` by setting the value to `true` in `hbase-site.xml`.

Apache Hadoop

Hadoop shell commands which reference the root directory ("/") do not work

Bug: [HDFS-5888](#)

Workaround: None

ResourceManager High Availability with manual failover does not work on secure clusters

Bug: [YARN-1640](#)

Workaround: Enable automatic failover; this requires ZooKeeper.

Apache Hive

*Select * fails on Parquet tables with the map data type*

Bug: [HIVE-6575](#)

Severity: Low

Workaround: Use the map's column name in the `SELECT` statement.

Cloudera Search

Malicious users could update information by circumventing Sentry checks

A sophisticated malicious user could update restricted content by setting the `update.distrib` parameter to bypass Sentry's index-level checks.

With Search for CDH 5.2 and later, Sentry always checks for index-level access control settings, preventing unauthorized updates.

Bug: None.

Severity: Medium.

Workaround: None.

Kerberos name rules were not followed

`SOLR_AUTHENTICATION_KERBEROS_NAME_RULES`, which is specified in `/etc/default/solr` or `/opt/cloudera/parcels/CDH-*/etc/default/solr` would sometimes not be respected, even if those rules were also specified using the `hadoop.security.auth_to_local` property in `SOLR_HDFS_CONFIG/core-site.xml`.

With Search for CDH 5.2 and later, Kerberos name rules are followed.

Bug: None.

Workaround: None.

Issues Fixed in CDH 5.1.x

The following topics describe issues fixed in CDH 5.1.x, from newest to oldest release. You can also review [What's New in CDH 5.1.x](#) on page 1854 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.1.7

Upstream Issues Fixed

CDH 5.1.7 includes the following issues fixed upstream:

- [HDFS-7960](#) - The full block report should prune zombie storages even if they're not empty
- [HDFS-7278](#) - Add a command that allows sysadmins to manually trigger full block reports from a DN
- [HDFS-6831](#) - Inconsistency between 'hdfs dfsadmin' and 'hdfs dfsadmin -help'
- [HDFS-7596](#) - NameNode should prune dead storages from storageMap
- [HDFS-7208](#) - NN doesn't schedule replication when a DN storage fails
- [HDFS-7575](#) - Upgrade should generate a unique storage ID for each volume
- [HDFS-6529](#) - Trace logging for RemoteBlockReader2 to identify remote datanode and file being read
- [YARN-570](#) - Time strings are formatted in different timezone
- [YARN-2251](#) - Avoid negative elapsed time in JHS/MRAM web UI and services
- [YARN-2588](#) - Standby RM does not transitionToActive if previous transitionToActive is failed with ZK exception.
- [HIVE-8634](#) - HiveServer2 fair scheduler queue mapping doesn't handle the secondary groups rules correctly
- [HIVE-8634](#) - HiveServer2 fair scheduler queue mapping doesn't handle the secondary groups rules correctly
- [HIVE-6403](#) - uncorrelated subquery is failing with `auto.convert.join=true`
- [HIVE-5945](#) - `ql.plan.ConditionalResolverCommonJoin.resolveMapJoinTask` also sums those tables which are not used in the child of this conditional task.
- [HIVE-8916](#) - Handle `user@domain` username under LDAP authentication
- [HIVE-8874](#) - Error Accessing HBase from Hive via Oozie on Kerberos 5.0.1 cluster
- [HIVE-9716](#) - Map job fails when table's LOCATION does not have scheme
- [HIVE-8784](#) - Querying partition does not work with JDO enabled against PostgreSQL
- [HUE-2484](#) - [beeswax] Configure support for Hive Server2 LDAP authentication
- [HUE-2446](#) - Migrating from CDH 4.7 to CDH 5.0.1+/Hue 3.5+ will fail
- [PARQUET-107](#) - Add option to disable summary metadata aggregation after MR jobs

- [SOLR-6268](#) - HdfsUpdateLog has a race condition that can expose a closed HDFS FileSystem instance and should close it's FileSystem instance if either inherited close method is called.
- [SOLR-6393](#) - Improve transaction log replay speed on HDFS.
- [SOLR-6403](#) - TransactionLog replay status logging.
- [IMPALA-1801](#) - external-data-source-executor leaking global jni refs
- [IMPALA-1794](#) - Fix infinite loop opening/closing file w/ invalid metadata
- [IMPALA-1674](#) - Fix serious memory leak in TSaslTransport
- [IMPALA-1668](#) - Fix leak of transport objects in TSaslServerTransport::Factory
- [IMPALA-1556](#) - TSaslTransport.read() should return available data before next frame
- [IMPALA-1565](#) - Python sasl client transport perf issue
- [IMPALA-1556](#) - Sasl transport should be wrapped with buffered transport
- [IMPALA-1442](#) - Better fix for non-buffered SASL transports The Thrift SASL implementation relies on the

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera (“Cloudera Products”), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled “Arbitrary remote code execution with InvokerTransformer” and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Issues Fixed in CDH 5.1.5

This is a maintenance release that fixes the following issues:

- [HDFS-7960](#) - The full block report should prune zombie storages even if they're not empty
- [HDFS-7278](#) - Add a command that allows sysadmins to manually trigger full block reports from a DN
- [HDFS-6831](#) - Inconsistency between 'hdfs dfsadmin' and 'hdfs dfsadmin -help'
- [HDFS-7596](#) - NameNode should prune dead storages from storageMap
- [HDFS-7208](#) - NN doesn't schedule replication when a DN storage fails
- [HDFS-7575](#) - Upgrade should generate a unique storage ID for each volume
- [HDFS-6529](#) - Trace logging for RemoteBlockReader2 to identify remote datanode and file being read
- [YARN-570](#) - Time strings are formatted in different timezone
- [YARN-2251](#) - Avoid negative elapsed time in JHS/MRAM web UI and services
- [YARN-2588](#) - Standby RM does not transitionToActive if previous transitionToActive is failed with ZK exception.
- [HIVE-8634](#) - HiveServer2 fair scheduler queue mapping doesn't handle the secondary groups rules correctly
- [HIVE-8634](#) - HiveServer2 fair scheduler queue mapping doesn't handle the secondary groups rules correctly

- [HIVE-6403](#) - uncorrelated subquery is failing with auto.convert.join=true
- [HIVE-5945](#) - ql.plan.ConditionalResolverCommonJoin.resolveMapJoinTask also sums those tables which are not used in the child of this conditional task.
- [HIVE-8916](#) - Handle user@domain username under LDAP authentication
- [HIVE-8874](#) - Error Accessing HBase from Hive via Oozie on Kerberos 5.0.1 cluster
- [HIVE-9716](#) - Map job fails when table's LOCATION does not have scheme
- [HIVE-8784](#) - Querying partition does not work with JDO enabled against PostgreSQL
- [HUE-2484](#) - [beeswax] Configure support for Hive Server2 LDAP authentication
- [HUE-2446](#) - Migrating from CDH 4.7 to CDH 5.0.1+/Hue 3.5+ will fail
- [PARQUET-107](#) - Add option to disable summary metadata aggregation after MR jobs
- [SOLR-6268](#) - HdfsUpdateLog has a race condition that can expose a closed HDFS FileSystem instance and should close it's FileSystem instance if either inherited close method is called.
- [SOLR-6393](#) - Improve transaction log replay speed on HDFS.
- [SOLR-6403](#) - TransactionLog replay status logging.
- [IMPALA-1801](#) - external-data-source-executor leaking global jni refs
- [IMPALA-1794](#) - Fix infinite loop opening/closing file w/ invalid metadata
- [IMPALA-1674](#) - Fix serious memory leak in TSaslTransport
- [IMPALA-1668](#) - Fix leak of transport objects in TSaslServerTransport::Factory
- [IMPALA-1556](#) - TSaslTransport.read() should return available data before next frame
- [IMPALA-1565](#) - Python sasl client transport perf issue
- [IMPALA-1556](#) - Sasl transport should be wrapped with buffered transport
- [IMPALA-1442](#) - Better fix for non-buffered SASL transports The Thrift SASL implementation relies on the

Issues Fixed in CDH 5.1.4

CDH 5.1.4 fixes the following issues, organized by component. See [What's New in CDH 5.1.x](#) on page 1854 for a list of the most important upstream problems fixed in this release.

HTTPS does not work on the HTTPS configured port

If you enable HTTPS (TLS/SSL) for YARN services, these services (including ResourceManager, NodeManager, and Job History Server) will not continue to use non-secure HTTP, but HTTPS does not work on the HTTPS configured port.

Bug: [YARN-1553](#)

Workaround: None.

Upstream Issues Fixed

In addition to the above, CDH 5.1.4 includes the following issues fixed upstream:

- [DATAFU-68](#) - SampleByKey can throw NullPointerException
- [HADOOP-11243](#) - SSLFactory shouldn't allow SSLv3
- [HADOOP-11156](#) - DelegateToFileSystem should implement getFsStatus(final Path f).
- [HDFS-7391](#) - Reenable SSLv2Hello in HttpFS
- [HDFS-7235](#) - DataNode#transferBlock should report blocks that don't exist using reportBadBlock
- [HDFS-7274](#) - Disable SSLv3 in HttpFS
- [HDFS-7005](#) - DFS input streams do not timeout
- [HDFS-6376](#) - Distcp data between two HA clusters requires another configuration
- [HDFS-6621](#) - Hadoop Balancer prematurely exits iterations
- [YARN-2273](#) - NPE in ContinuousScheduling thread when we lose a node
- [YARN-2566](#) - DefaultContainerExecutor should pick a working directory randomly
- [YARN-2588](#) - Standby RM does not transitionToActive if previous transitionToActive is failed with ZK exception.
- [YARN-2641](#) - Decommission nodes on -refreshNodes instead of next NM-RM heartbeat
- [YARN-2608](#) - FairScheduler: Potential deadlocks in loading alloc files and clock access
- [HBASE-12376](#) - HBaseAdmin leaks ZK connections if failure starting watchers (ConnectionLossException)

- [HBASE-12366](#) - Add login code to HBase Canary tool
- [HBASE-12098](#) - User granted namespace table create permissions can't create a table
- [HBASE-12087](#) - [0.98] Changing the default setting of hbase.security.access.early_out to true
- [HBASE-11896](#) - LoadIncrementalHFiles fails in secure mode if the namespace is specified
- [HBASE-12054](#) - bad state after NamespaceUpgrade with reserved table names
- [HBASE-12460](#) - Moving Chore to hbase-common module
- [HIVE-5643](#) - ZooKeeperHiveLockManager.getQuorumServers incorrectly appends the custom zk port to quorum hosts
- [HIVE-8675](#) - Increase thrift server protocol test coverage
- [HIVE-8827](#) - Remove SSLv2Hello from list of disabled protocols
- [HIVE-8182](#) - beeline fails when executing multiple-line queries with trailing spaces
- [HIVE-8330](#) - HiveResultSet.findColumn() parameters are case sensitive
- [HIVE-5994](#) - ORC RLEv2 encodes wrongly for large negative BIGINTs (64 bits)
- [HIVE-7629](#) - Problem in SMB Joins between two Parquet tables
- [HIVE-6670](#) - ClassNotFoundException with Serde
- [HIVE-6409](#) - FileOutputCommitterContainer::commitJob() cancels delegation tokens too early.
- [HIVE-7647](#) - Beeline does not honor --headerInterval and --color when executing with \
- [HIVE-7441](#) - Custom partition scheme gets rewritten with hive scheme upon concatenate
- [HIVE-5871](#) - Use multiple-characters as field delimiter
- [HIVE-1363](#) - SHOW TABLE EXTENDED LIKE command does not strip single/double quotes
- [HIVE-5989](#) - Hive metastore authorization check is not threadsafe
- [HUE-2438](#) - [core] Disable SSLv3 for Poodle vulnerability
- [HUE-2291](#) - [oozie] Faster dashboard display
- [IMPALA-1334](#) - Impala does not map principals to lowercase, affecting Sentry authorisation
- [IMPALA-1251](#) - High-offset queries hang
- [IMPALA-1338](#) - HDFS does not return all ACLs in getAclStatus()
- [IMPALA-1279](#) - Impala does not employ ACLs when checking path permissions for LOAD and INSERT
- [OOZIE-2034](#) - Disable SSLv3 (POODLEbleed vulnerability)
- [OOZIE-2063](#) - Cron syntax creates duplicate actions
- [SENTRY-428](#) - Sentry service should periodically renew the server kerberos ticket
- [SENTRY-431](#) - Sentry db provider client should attempt to refresh kerberos ticket before connection
- [SPARK-3606](#) - Spark-on-Yarn AmlpFilter does not work with Yarn HA

Issues Fixed in CDH 5.1.3

CDH 5.1.3 fixes the following issues, organized by component. See [New Features and Changes in CDH 5](#) on page 1812 for a list of the most important upstream problems fixed in this release.

Apache Hadoop

The default setting of `dfs.client.block.write.replace-datanode-on-failure.policy` can cause an unrecoverable error in small clusters

The default setting of `dfs.client.block.write.replace-datanode-on-failure.policy` (DEFAULT) can cause an unrecoverable error in a small cluster during HBase rolling restart.

Bug: [HDFS-4257](#)

Workaround: Set `dfs.client.block.write.replace-datanode-on-failure.policy` to NEVER for 1- 2- or 3-node clusters, and leave it as DEFAULT for all other clusters. Leave `dfs.client.block.write.replace-datanode-on-failure.enable` set to true.

Upstream Issues Fixed

In addition to the above, CDH 5.1.3 includes the following issues fixed upstream.

- [HADOOP-11035](#) - distcp on mr1(branch-1) fails with NPE using a short relative source path.

- [HBASE-10188](#) - Hide ServerName constructor
- [HBASE-10012](#) - Hide ServerName constructor
- [HBASE-11349](#) - [Thrift] support authentication/impersonation
- [HBASE-11446](#) - Reduce the frequency of RNG calls in SecureWALCellCodec#EncryptedKvEncoder
- [HBASE-11457](#) - Increment HFile block encoding IVs accounting for cipher's internal use
- [HBASE-11474](#) - [Thrift2] support authentication/impersonation
- [HBASE-11565](#) - Stale connection could stay for a while
- [HBASE-11627](#) - RegionSplitter's rollingSplit terminated with "/ by zero", and the `_balancedSplit` file was not deleted properly
- [HBASE-11788](#) - hbase is not deleting the cell when a Put with a KeyValue, KeyValue.Type.Delete is submitted
- [HBASE-11828](#) - Callers of ServerName.valueOf should use equals and not ==
- [HDFS-4257](#) - The ReplaceDatanodeOnFailure policies could have a forgiving option
- [HDFS-6776](#) - Using distcp to copy data between insecure and secure cluster via webhdfs does not work
- [HDFS-6908](#) - Incorrect snapshot directory diff generated by snapshot deletion
- [HUE-2247](#) - [Impala] Support pass-through LDAP authentication
- [HUE-2295](#) - [librdbs] External oracle DB connection is broken due to a typo
- [HUE-2273](#) - [desktop] Blacklisting apps with existing document will break home page
- [HUE-2318](#) - [desktop] Documents shared with write group permissions are not editable
- [HIVE-5087](#) - Rename npath UDF to matchpath
- [HIVE-6820](#) - HiveServer(2) ignores HIVE_OPTS
- [HIVE-7635](#) - Query having same aggregate functions but different case throws IndexOutOfBoundsException
- [IMPALA-958](#) - Excessively long query plan serialization time in FE when querying huge tables
- [IMPALA-1091](#) - Improve TScanRangeLocation struct and associated code
- [OOZIE-1989](#) - NPE during a rerun with forks
- [YARN-1458](#) - FairScheduler: Zero weight can lead to livelock

Issues Fixed in CDH 5.1.2

CDH 5.1.2 fixes the following issues, organized by component. See [What's New in CDH 5.1.x](#) on page 1854 for a list of the most important upstream problems fixed in this release.

Apache Hadoop

Jobs can hang on NodeManager decommission owing to a race condition when continuous scheduling is enabled.

Bug: [YARN-2273](#)

Workaround: Disable continuous scheduling by setting `yarn.scheduler.fair.continuous-scheduling-enabled` to false

Apache HBase

Sending a large amount of invalid data to the Thrift service can cause it to crash

Bug: [HBASE-11052](#).

Workaround: None. This is a longstanding problem, not a new issue in CDH 5.1.

The metric `ageOfLastShippedOp` never decreases

This can cause it to appear as though the cluster is in an inconsistent state even when there is no problem.

Bug: [HBASE-11143](#).

Workaround: None.

Upstream Issues Fixed

In addition to the above, CDH 5.1.2 includes the following issues fixed upstream.

- [FLUME-2438](#)
- [HBASE-11052](#)

- [HBASE-11143](#)
- [HBASE-11609](#)
- [HDFS-6114](#)
- [HDFS-6640](#)
- [HDFS-6703](#)
- [HDFS-6788](#)
- [HDFS-6825](#)
- [HUE-2211](#)
- [HUE-2223](#)
- [HUE-2232](#)
- [HIVE-5515](#)
- [HIVE-6495](#)
- [HIVE-7450](#)
- [IMPALA-1093](#)
- [IMPALA-1107](#)
- [IMPALA-1131](#)
- [IMPALA-1142](#)
- [IMPALA-1149](#)
- [MAPREDUCE-5966](#)
- [MAPREDUCE-5979](#)
- [MAPREDUCE-6012](#)
- [OOZIE-1920](#)
- [PARQUET-19](#)
- [SENTRY-363](#)
- [YARN-2273](#)
- [YARN-2274](#)
- [YARN-2313](#)
- [YARN-2352](#)
- [YARN-2359](#)

Issues Fixed in CDH 5.1.0

Apache Hadoop

HDFS

The same DataNodes may appear in the NameNode web UI in both the live and dead node lists

Bug: [HDFS-6180](#)

Workaround: None

MapReduce

YARN Fair Scheduler's Cluster Utilization Threshold check is broken

Bug: [YARN-1640](#)

Workaround: Set the `yarn.scheduler.fair.preemption.cluster-utilization-threshold` property in `yarn-site.xml` to -1.

ResourceManager High Availability with manual failover does not work on secure clusters

Bug: [YARN-2155](#)

Workaround: Enable automatic failover; this requires ZooKeeper.

Apache HBase

MapReduce over HBase Snapshot bypasses HBase-level security

The MapReduce over HBase Snapshot bypasses HBase-level security completely since the files are read from the HDFS directly. The user who is running the scan/job has to have read permissions to the data and snapshot files.

Bug: [HBASE-8369](#)

Workaround: MapReduce users must be trusted to process/view all data in HBase.

HBase snapshots now saved to the /<hbase>/hbase-snapshot directory

HBase snapshots are now saved to the /<hbase>/hbase-snapshot directory instead of the /hbase/.snapshot directory. This was a conflict introduced by the HDFS snapshot feature in Hadoop 2.2/CDH 5 HDFS.

Bug: [HBASE-8352](#)

Workaround: This should be handled in the upgrade process.

Hue

Oozie jobs don't support ResourceManager HA in YARN

If the ResourceManager fails, the workflow will fail.

Bug: None

Severity: Medium

Workaround: None

Apache Oozie

Oozie HA does not work properly with HCatalog integration or SLA notifications

This issue appears when you are using HCatalog as a data dependency in a coordinator; using HCatalog from an action (for example, Pig) works correctly.

Bug: [OOZIE-1492](#)

Workaround: None

Issues Fixed in CDH 5.0.x

The following topics describe issues fixed in CDH 5.0.x, from newest to oldest release. You can also review [What's New in CDH 5.0.x](#) on page 1858 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5.0.6

Upstream Issues Fixed

The following upstream issues are fixed in CDH 5.0.6:

- [HDFS-7960](#) - The full block report should prune zombie storages even if they're not empty
- [HDFS-7278](#) - Add a command that allows sysadmins to manually trigger full block reports from a DN
- [HDFS-6831](#) - Inconsistency between `hdfs dfsadmin` and `hdfs dfsadmin -help`
- [HDFS-7596](#) - NameNode should prune dead storages from `storageMap`
- [HDFS-7208](#) - NN does not schedule replication when a DN storage fails
- [HDFS-7575](#) - Upgrade should generate a unique storage ID for each volume
- [YARN-570](#) - Time strings are formatted in different timezone
- [YARN-2251](#) - Avoid negative elapsed time in JHS/MRAM web UI and services
- [HIVE-8874](#) - Error Accessing HBase from Hive via Oozie on Kerberos 5.0.1 cluster
- [SOLR-6268](#) - `HdfsUpdateLog` has a race condition that can expose a closed HDFS `FileSystem` instance and should close its `FileSystem` instance if either inherited close method is called.
- [SOLR-6393](#) - Improve transaction log replay speed on HDFS.
- [SOLR-6403](#) - `TransactionLog` replay status logging.

Issues Fixed in CDH 5.0.5

"POODLE" Vulnerability on TLS/SSL enabled ports

The POODLE (Padding Oracle On Downgraded Legacy Encryption) attack takes advantage of a cryptographic flaw in the obsolete SSLv3 protocol, after first forcing the use of that protocol. The only solution is to disable SSLv3 entirely.

This requires changes across a wide variety of components of CDH and Cloudera Manager in all current versions. CDH 5.0.5 provides these changes for CDH 5.0.x deployments.

For more information, see the [Cloudera Security Bulletin](#).

Apache Hadoop Distributed Cache Vulnerability

The Distributed Cache Vulnerability allows a malicious cluster user to expose private files owned by the user running the YARN NodeManager process. For more information, see the [Cloudera Security Bulletin](#).

Upstream Issues Fixed

CDH 5.0.4 includes the following issues fixed upstream.

- [HADOOP-11243](#) - SSLFactory shouldn't allow SSLv3
- [HDFS-7274](#) - Disable SSLv3 in HttpFS
- [HDFS-7391](#) - Reenable SSLv2Hello in HttpFS
- [HBASE-12376](#) - HBaseAdmin leaks ZK connections if failure starting watchers (ConnectionLossException)
- [HIVE-8675](#) - Increase thrift server protocol test coverage
- [HIVE-8827](#) - Remove SSLv2Hello from list of disabled protocols protocols
- [HUE-2438](#) - [core] Disable SSLv3 for Poodle vulnerability
- [OOZIE-2034](#) - Disable SSLv3 (POODLEbleed vulnerability)
- [OOZIE-2063](#) - Cron syntax creates duplicate actions

Issues Fixed in CDH 5.0.4

Upstream Issues Fixed

CDH 5.0.4 includes the following issues fixed upstream.

- [FLUME-2438](#)
- [HBASE-11609](#)
- [HDFS-6044](#)
- [HDFS-6529](#)
- [HDFS-6618](#)
- [HDFS-6622](#)
- [HDFS-6640](#)
- [HDFS-6647](#)
- [HDFS-6703](#)
- [HDFS-6788](#)
- [HIVE-5515](#)
- [HIVE-7459](#)
- [HUE-2166](#)
- [HUE-2249](#)
- [IMPALA-1019](#)
- [MAPREDUCE-5966](#)
- [MAPREDUCE-5979](#)
- [OOZIE-1920](#)
- [PARQUET-19](#)
- [SPARK-1930](#)
- [YARN-1550](#)
- [YARN-2061](#)
- [YARN-2132](#)

Issues Fixed in CDH 5.0.3

The following topics describe known issues fixed in CDH 5.0.3. See [What's New in CDH 5.0.x](#) on page 1858 for a list of the most important upstream problems fixed in this release.

CDH 5 Release Notes

Apache Hadoop

MapReduce

YARN Fair Scheduler's Cluster Utilization Threshold check is broken

Bug: [YARN-2155](#)

Workaround: Set the `yarn.scheduler.fair.preemption.cluster-utilization-threshold` property in `yarn-site.xml` to -1.

Apache Oozie

When Oozie is configured to use MRv1 and TLS/SSL, YARN / MRv2 libraries are erroneously included in the classpath instead

This problem causes much of the configured Oozie functionality to be unusable.

Bug: None

Workaround: Use a different configuration (non-TLS/SSL or YARN), if possible.

Upstream Issues Fixed

In addition to the above, CDH 5.0.2 includes the following issues fixed upstream.

- [FLUME-2245](#)
- [FLUME-2416](#)
- [HBASE-10871](#)
- [HDFS-5891](#)
- [HDFS-6021](#)
- [HDFS-6077](#)
- [HDFS-6340](#)
- [HDFS-6475](#)
- [HDFS-6510](#)
- [HDFS-6527](#)
- [HDFS-6563](#)
- [HUE-1928](#)
- [HUE-2184](#)
- [HUE-2085](#)
- [HUE-2192](#)
- [HUE-2193](#)
- [OOZIE-1621](#)
- [OOZIE-1890](#)
- [OOZIE-1907](#)
- [SOLR-5593](#)
- [SOLR-5915](#)
- [SOLR-6161](#)
- [YARN-1550](#)
- [YARN-2155](#)

Issues Fixed in CDH 5.0.2

The following topics describe known issues fixed in CDH 5.0.2. See [What's New in CDH 5.0.x](#) on page 1858 for a list of the most important upstream problems fixed in this release.

Apache Hadoop

CDH 5 clients running releases 5.0.1 and earlier cannot use WebHDFS to connect to a CDH 4 cluster

For example, a `hadoop fs -ls webhdfs` command run from the CDH 5 client to the CDH 4 cluster produces an error such as the following:

```
Found 21 items
ls: Invalid value for webhdfs parameter "op": No enum const class
org.apache.hadoop.hdfs.web.resources.GetOpParam.Op.GETACLSTATUS
```

Bug: [HDFS-6326](#)

Workaround: None; note that this is fixed as of CDH 5.0.2.

Apache HBase

Endless Compaction Loop

If an empty HFile whose max timestamp is past its TTL (time-to-live) is selected for compaction, it is compacted into another empty HFile, which is selected for compaction, creating an endless compaction loop.

Bug: [HBASE-10371](#)

Workaround: None

Upstream Issues Fixed

In addition to the above, CDH 5.0.2 includes the following issues fixed upstream.

- [HADOOP-10556](#)
- [HADOOP-10638](#)
- [HADOOP-10639](#)
- [HADOOP-10658](#)
- [HBASE-6690](#)
- [HBASE-10312](#)
- [HBASE-10371](#)
- [HDFS-6326](#)
- [HIVE-5380](#)
- [HIVE-6913](#)
- [PIG-3677](#)
- [YARN-2073](#)

Issues Fixed in CDH 5.0.1

Apache Hadoop

HDFS

NameNode LeaseManager may crash

Bug: [HDFS-6148](#)/[HDFS-6094](#)

Workaround: Restart the NameNode.

Some group mapping providers can cause the NameNode to crash

In certain environments, some group mapping providers can cause the NameNode to segfault and crash.

Bug: [HADOOP-10442](#)

Workaround: Configure either `ShellBasedUnixGroupsMapping` in Hadoop or configure SSSD in the operating system on the NameNode.

Apache Hive

CREATE TABLE AS SELECT (CTAS) does not work with Parquet files

Since CTAS does not work with Parquet files, the following example will return null values.

```
CREATE TABLE test_data(column1 string);
LOAD DATA LOCAL INPATH './data.txt' OVERWRITE INTO TABLE test_data;

CREATE TABLE parquet_test
ROW FORMAT SERDE 'parquet.hive.serde.ParquetHiveSerDe'
  STORED AS
  INPUTFORMAT 'parquet.hive.DeprecatedParquetInputFormat'
  OUTPUTFORMAT 'parquet.hive.DeprecatedParquetOutputFormat'
AS
  SELECT column1 FROM test_data;

SELECT * FROM parquet_test;
SELECT column1 FROM parquet_test;
```

Bug: [HIVE-6375](#)

Workaround: A workaround for this is to follow up a CREATE TABLE query with an INSERT OVERWRITE TABLE SELECT * as in the example below.

```
CREATE TABLE parquet_test (column1 string)
ROW FORMAT SERDE 'parquet.hive.serde.ParquetHiveSerDe'
  STORED AS
  INPUTFORMAT 'parquet.hive.DeprecatedParquetInputFormat'
  OUTPUTFORMAT 'parquet.hive.DeprecatedParquetOutputFormat';
INSERT OVERWRITE TABLE parquet_test SELECT * from test_data;
```

Apache Oozie

The oozie-workflow-0.4.5 schema has been removed

Workflows using schema 0.4.5 will no longer be accepted by Oozie because this schema definition version has been removed.

Bug: [OOZIE-1768](#)

Workaround: Use schema 0.5. It's backwards compatible with 0.4.5, so updating the workflow is as simple as changing the schema version number.

Upstream Issues Fixed

In addition to the above, CDH 5.0.1 includes the following issues fixed upstream.

- [HADOOP-10442](#) - Group look-up can cause segmentation fault when a certain JNI-based mapping module is used.
- [HADOOP-10456](#) - Bug in Configuration.java exposed by Spark (ConcurrentModificationException)
- [HDFS-5064](#) - Standby checkpoints should not block concurrent readers
- [HDFS-6039](#) - Uploading a File under a Dir with default ACLs throws "Duplicated ACLFeature"
- [HDFS-6094](#) - The same block can be counted twice towards safe mode threshold
- [HDFS-6231](#) - DFSClient hangs infinitely if using hedged reads and all eligible DataNodes die
- [HIVE-6495](#) - TableDesc.getDeserializer() should use correct classloader when calling Class.forName()
- [HIVE-6575](#) - select * fails on parquet table with map data type
- [HIVE-6648](#) - Fixed permission inheritance for multi-partitioned tables
- [HIVE-6740](#) - Fixed addition of Avro JARs to classpath
- [HUE-2061](#) - Task logs are not retrieved if containers not on the same host
- [OOZIE-1794](#) - java-opts and java-opt in the Java action don't always work properly in YARN
- [SOLR-5608](#) - Frequently reproducible failures in CollectionsAPIDistributedZkTest#testDistribSearch
- [YARN-1924](#) - STATE_STORE_OP_FAILED happens when ZKRMStateStore tries to update app(attempt) before storing it

Issues Fixed in CDH 5.0.0

Apache Flume

AsyncHBaseSink does not work in CDH 5 Beta 1 and CDH 5 Beta 2

Bug: None

Workaround: Use the HBASE sink (`org.apache.flume.sink.hbase.HBaseSink`) to write to HBase in CDH 5 Beta releases.

Apache Hadoop

HDFS

DataNode can consume 100 percent of one CPU

A narrow race condition can cause one of the threads in the DataNode process to get stuck in a tight loop and consume 100 percent of one CPU.

Bug: [HDFS-5922](#)

Workaround: Restart the DataNode process.

HDFS NFS gateway does not work with Kerberos-enabled clusters

Bug: [HDFS-5898](#)

Workaround: None.

Cannot browse filesystem via NameNode Web UI if any directory has the sticky bit set

When listing any directory which contains an entry that has the sticky bit permission set, for example `/tmp` is often set this way, nothing will appear where the list of files or directories should be.

Bug: [HDFS-5921](#)

Workaround: Use the Hue File Browser.

Appending to a file that has been snapshotted previously will append to the snapshotted file as well

If you append content to a file that exists in snapshot, the file in snapshot will have the same content appended to it, invalidating the original snapshot.

Bug: See also [HDFS-5343](#)

Workaround: None

MapReduce

In MRv2 (YARN), the JobHistory Server has no information about a job if the ApplicationMasters fails while the job is running

Bug: None

Workaround: None.

Apache HBase

An empty rowkey is treated as the first row of a table

An empty rowkey is allowed in HBase, but it was treated as the first row of the table, even if it was not in fact the first row. Also, multiple rows with empty rowkeys caused issues.

Bug: [HBASE-3170](#)

Workaround: Do not use empty rowkeys.

Apache Hive

Hive queries that combine multiple splits and query large tables fail on YARN

Hive queries that scan large tables, or perform map side joins may fail with the following exception when the query is run using YARN:

```
java.io.IOException: Max block location exceeded for split:
InputFormatClass: org.apache.hadoop.mapred.TextInputFormat
```

```

splitsize: 21 maxsize: 10
at
org.apache.hadoop.mapreduce.split.JobSplitWriter.writeOldSplits(JobSplitWriter.java:162)
at
org.apache.hadoop.mapreduce.split.JobSplitWriter.createSplitFiles(JobSplitWriter.java:87)
at org.apache.hadoop.mapreduce.JobSubmitter.writeOldSplits(JobSubmitter.java:540)
at org.apache.hadoop.mapreduce.JobSubmitter.writeSplits(JobSubmitter.java:510)
at org.apache.hadoop.mapreduce.JobSubmitter.submitJobInternal(JobSubmitter.java:392)
at org.apache.hadoop.mapreduce.Job$10.run(Job.java:1268)
at org.apache.hadoop.mapreduce.Job$10.run(Job.java:1265)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:415)
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1491)
at org.apache.hadoop.mapreduce.Job.submit(Job.java:1265)
at org.apache.hadoop.mapred.JobClient$1.run(JobClient.java:562)
at org.apache.hadoop.mapred.JobClient$1.run(JobClient.java:557)
at java.security.AccessController.doPrivileged(Native Method)
at javax.security.auth.Subject.doAs(Subject.java:415)
at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1491)
at org.apache.hadoop.mapred.JobClient.submitJobInternal(JobClient.java:557)
at org.apache.hadoop.mapred.JobClient.submitJob(JobClient.java:548)
at org.apache.hadoop.hive.ql.exec.mr.ExecDriver.execute(ExecDriver.java:425)
at org.apache.hadoop.hive.ql.exec.mr.MapRedTask.execute(MapRedTask.java:136)
at org.apache.hadoop.hive.ql.exec.Task.executeTask(Task.java:151)
at org.apache.hadoop.hive.ql.exec.TaskRunner.runSequential(TaskRunner.java:65)

```

Bug: [MAPREDUCE-5186](#)

Workaround: Set `mapreduce.job.max.split.locations` to a high value such as 100.

Files in Avro tables no longer have .avro extension

As of CDH 4.3.0 Hive no longer creates files in Avro tables with the `.avro` extension by default. This does not cause any problems in Hive, but could affect downstream components such as Pig, MapReduce, or Sqoop 1 that expect files with the `.avro` extension.

Bug: None

Workaround: Manually set the extension to `.avro` before using a query that inserts data into your Avro table. Use the following `set` statement:

```
set hive.output.file.extension=".avro";
```

Apache Oozie

Oozie does not work seamlessly with ResourceManager HA

Oozie workflows are not recovered on ResourceManager failover when ResourceManager HA is enabled. Further, users cannot specify the `clusterId` for JobTracker to work against either ResourceManager.

Bug: None

Workaround: On non-secure clusters, users are required to specify either of the ResourceManagers' `host:port`. For secure clusters, users are required to specify the Active ResourceManager's `host:port`.

When using Oozie HA with security enabled, some znodes have world ACLs

Oozie High Availability with security enabled will still work, but a malicious user or program can alter znodes used by Oozie for locking, possibly causing Oozie to be unable to finish processing certain jobs.

Bug: [OOZIE-1608](#)

Workaround: None

Oozie and Sqoop 2 may need additional configuration to work with YARN

In CDH 5, MRv2 (YARN) MapReduce 2.0 is recommended over the Hadoop 0.20-based MRv1. The default configuration may not reflect this in Oozie and Sqoop 2 in CDH 5 Beta 2, however, unless you are using Cloudera Manager.

Bug: None

Workaround: Check the value of `CATALINA_BASE` in `/etc/oozie/conf/oozie-env.sh` (if you are running an Oozie server) and `/etc/default/sqoop2-server` (if you are using a Sqoop 2 server). You should also ensure that `CATALINA_BASE` is correctly set in your environment if you are invoking `/usr/bin/sqoop2-server` directly instead of using the service init scripts. For Oozie, `CATALINA_BASE` should be set to `/usr/lib/oozie/oozie-server` for YARN, or `/usr/lib/oozie/oozie-server-0.20` for MRv1. For Sqoop 2, `CATALINA_BASE` should be set to `/usr/lib/sqoop2/sqoop-server` for YARN, or `/usr/lib/sqoop2/sqoop-server-0.20` on MRv1.

Cloudera Search

Creating cores using the web UI with default values causes the system to become unresponsive

You can use the Solr Server web UI to create new cores. If you click **Create Core** without making any changes to the default attributes, the server may become unresponsive. Checking the log for the server shows a repeated error that begins:

```
ERROR org.apache.solr.cloud.Overseer: Exception in Overseer main queue loop
java.lang.IllegalArgumentException: Path must not end with / character
```

Bug: Solr-5813

Workaround: To avoid this issue, do not create cores without first updating values for the new core in the web UI. For example, you might enter a new name for the core to be created.

If you created a core with default settings and are seeing this error, you can address the problem by finding which node is having problems and removing that node. Find the problematic node by using a tool that can inspect ZooKeeper, such as the Solr Admin UI. Using such a tool, examine items in the ZooKeeper queue, reviewing the properties for the item. The problematic node will have an item in its queue with the property `collection=""`.

Remove the node with the item with the `collection=""` property using a ZooKeeper management tool. For example, you can remove nodes using the ZooKeeper command line tool or recent versions of HUE.

Issues Fixed in CDH 5 Beta Releases

The following topics describe issues fixed in CDH 5 Beta 2, after being discovered in CDH 5 Beta 1. You can also review [What's New In CDH 5 Beta Releases](#) on page 1861 or [Known Issues in CDH 5](#) on page 1944.

Issues Fixed in CDH 5 Beta 2

Apache Hadoop

MapReduce

ResourceManager High Availability does not work on secure clusters

If JobTrackers in an High Availability configuration are shut down, migrated to new hosts, then restarted, no JobTracker becomes active. The logs show a `Mismatched address` exception.

Bug: None

Workaround: None.

Default port conflicts

By default, the Shuffle Handler (which runs inside the YARN NodeManager), the REST server, and many third-party applications, all use port 8080. This will result in conflicts if you deploy more than one of them without reconfiguring the default port.

Bug: None

Workaround: Make sure at most one service uses port 8080. To reconfigure the REST server, follow [these instructions](#). To change the default port for the Shuffle Handler, set the value of `mapreduce.shuffle.port` in `mapred-site.xml` to an unused port.

JobTracker memory leak

The JobTracker has a memory leak caused by subtleties in the way `UserGroupInformation` interacts with the file-system cache. The number of cached file system objects can grow without bound.

Bug: [MAPREDUCE-5508](#)

Workaround: Set `keep.failed.task.files` to true, which will sidestep the memory leak but require job staging directories to be cleaned out manually.

Hue

Running a Hive Beeswax metastore on the same host as the Hue server will result in Simple Authentication and Security Layer (SASL) authentication failures on a Kerberos-enabled cluster

Bug: None

Workaround: The simple workaround is to run the metastore server remotely on a different host and make sure all Hive and Hue configurations properly refer to it. A more complex workaround is to adjust network configurations to ensure that reverse DNS properly resolves the host's address to its fully qualified-domain name (FQDN) rather than `localhost`.

The Pig shell does not work when NameNode uses a wildcard address

The Pig shell does not work from Hue if you use a wildcard for the NameNode's RPC or HTTP bind address. For example, `dfs.namenode.http-address` must be a real, routable address and port, not `0.0.0.0.<port>`.

Bug: [HUE-1060](#)

Workaround: Use a real, routable address and port, not `0.0.0.0.<port>`, for the NameNode; or use the Pig application directly, rather than from Hue.

Apache Sqoop

Oozie and Sqoop 2 may need additional configuration to work with YARN

In CDH 5, MRv2 (YARN) MapReduce 2.0 is recommended over the Hadoop 0.20-based MRv1. The default configuration may not reflect this in Oozie and Sqoop 2 in CDH 5 Beta 2, however, unless you are using Cloudera Manager.

Bug: None

Workaround: Check the value of `CATALINA_BASE` in `/etc/oozie/conf/oozie-env.sh` (if you are running an Oozie server) and `/etc/default/sqoop2-server` (if you are using a Sqoop 2 server). You should also ensure that `CATALINA_BASE` is correctly set in your environment if you are invoking `/usr/bin/sqoop2-server` directly instead of using the service `init` scripts. For Oozie, `CATALINA_BASE` should be set to `/usr/lib/oozie/oozie-server` for YARN, or `/usr/lib/oozie/oozie-server-0.20` for MRv1. For Sqoop 2, `CATALINA_BASE` should be set to `/usr/lib/sqoop2/sqoop-server` for YARN, or `/usr/lib/sqoop2/sqoop-server-0.20` on MRv1.

Apache Sentry (incubating)

Sentry allows unauthorized access to a directory whose name includes the scratch directory name as a prefix

As an example, if the scratch directory path is `/tmp/hive`, and you create a directory `/tmp/hive-data`, Sentry allows unauthorized read/write access to `/tmp/hive-data`.

Bug: None

Workaround: For external tables or data export location, do not use a pathname that includes the scratch directory name as a prefix. For example, if the scratch directory is `/tmp/hive`, do not locate external tables or exported data in `/tmp/hive-data` or any directory whose path uses `"/tmp/hive-"` as a prefix.

Apache Oozie

Oozie Hive action against HiveServer2 fails on a secure cluster

Workaround: None

Fixed Issues in Apache Impala

The following sections describe the major issues fixed in each Impala release.

For known issues that are currently unresolved, see [Apache Impala Known Issues](#) on page 1995.

Issues Fixed in Impala for CDH 5.16.2

For the full list of fixed issues for all CDH components in CDH 5.16.2, see [Issues Fixed in CDH 5.16.x](#) on page 2054. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-6323](#) - Impala now supports a constant in the window specifications.
- [IMPALA-7960](#) - Impala now returns a correct result when comparing `TIMESTAMP` to a string literal in a binary predicate where the `TIMESTAMP` is casted to `VARCHAR` of smaller length.
- [IMPALA-7961](#) - Fixed an issue where queries running with the `SYNC_DDL` query option can fail when the Catalog Server is under a heavy load with concurrent catalog operations of long-running DDLs.
- [IMPALA-8058](#) - Fixed cardinality estimates for HBase queries, which could sometimes yield hugely high numbers.
- [IMPALA-8109](#) - Impala can now read the gzip files bigger than 2 GB.
- [IMPALA-8212](#) - Fixed a race condition in the Kerberos authentication code.

Issues Fixed in Impala for CDH 5.16.1

For the full list of fixed issues for all CDH components in CDH 5.16.1, see [Issues Fixed in CDH 5.16.x](#) on page 2054. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-6086](#) - Require the `SELECT` privilege on the database for built-in function calls.
- [IMPALA-6451](#) - Fixed the `AuthorizationException` in CTAS for Kudu tables.
- [IMPALA-6479](#) - `DESCRIBE` now respects column level privileges and only shows the columns that the user has the privilege to view.
- [IMPALA-6571](#) - Fixed the `NullPointerException` in `SHOW CREATE TABLE` for HBase tables.
- [IMPALA-7225](#) - `REFRESH..PARTITION` no longer reset the number of rows in a partition.
- [IMPALA-7272](#) - Fixed the crash in `StringMinMaxFilter`.
- [IMPALA-7360](#) - Fixed an issue where Avro scanner sometimes skipped blocks when skip marker was on HDFS block boundary.
- [IMPALA-7419](#) - Fixed the `NullPointerException` in `SimplifyConditionalsRule`.
- [IMPALA-7483](#) - `impalad/catalogd` on JVM deadlock now get aborted.
- [IMPALA-7520](#) - Fixed the `NullPointerException` in `SentryProxy`.

Issues Fixed in Impala for CDH 5.15.2

For the full list of fixed issues for all CDH components in CDH 5.15.2, see [Issues Fixed in CDH 5.15.x](#) on page 2063. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-6907](#) - Now Impala correctly closes all stale connections to removed impala cluster members.
- [IMPALA-7225](#) - Fixed an issue where the `REFRESH...PARTITION` statement caused statistics for the refreshed partition to be automatically reset to `-1 (unknown)`. With the fix, statistics will be changed only if an explicit `COMPUTE STATS` statement is issued for an object.
- [IMPALA-7272](#) - Fixed a crash caused by a memory management problem when the query execution requires finding strings inside a range defined by the lesser-than and greater-than comparisons.
- [IMPALA-7360](#) - Fixed an issue where Impala could incorrectly skip data if a record separator in a sequence-based file (Avro, RC or sequence file) straddled an HDFS block boundary.
- [IMPALA-7537](#) - Fixed a security issue where `REVOKE ALL ON SERVER` did not have a permanent effect if the `ALL` permission was granted using the `WITH GRANT` option. Running `INVALIDATE METADATA` no longer causes the permission to reappear.
- [IMPALA-7585](#) - Fixed an issue in `KRPC`, which can cause slow or hung queries on non-secured clusters.

Issues Fixed in Impala for CDH 5.15.1

For the full list of fixed issues for all CDH components in CDH 5.15.1, see [Issues Fixed in CDH 5.15.x](#) on page 2063. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-6687](#) - Fixed `INSERT` with mixed case in partition column names.
- [IMPALA-6822](#) - Added a query option to control shuffling by distinct expressions.
- [IMPALA-6847](#) - Added a work around high memory estimates for admission control.
- [IMPALA-6908](#) - `IsConnResetTException()` should include `ECONNRESET`.
- [IMPALA-6934](#) - Corrected the wrong results with `EXISTS` subqueries that contain `ORDER BY`, `LIMIT`, and `OFFSET`.
- [IMPALA-7014](#) - Disabled the stacktrace symbolisation by default.

- [IMPALA-7078](#) - Reduced the queue size based on num_scanner_threads.
- [IMPALA-7078](#) - Improved memory consumption of wide Avro scans.
- [IMPALA-7288](#) - Fixed the Codegen crash in FinalizeModule.
- [IMPALA-7298](#) - Impala no longer passes IP address as hostname in Kerberos principal.

Issues Fixed in Impala for CDH 5.15.0

For the full list of fixed issues for all CDH components in CDH 5.15.0, see [Issues Fixed in CDH 5.15.x](#) on page 2063. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-4315](#) - Allow `USE` and `SHOW TABLES` if there is at least one table in a database where the user has table or column privileges.
- [IMPALA-4323](#) - The `SET ROW FORMAT` clause was added to the `ALTER TABLE` statement for the `TEXT` or `SEQUENCE` file formats.
- [IMPALA-4886](#) - Table metrics are available in the catalog web UI.
- [IMPALA-5654](#) - Disallows explicitly setting the Kudu table name property for managed Kudu tables in a `CREATE TABLE` and `ALTER TABLE` statements, e.g. `CREATE TABLE t (i INT) STORED AS KUDU TBLPROPERTIES('kudu.table_name' = 'some_name')`.
- [IMPALA-6549](#) - The file handle cache, controlled by the `max_cached_file_handles` flag, is enabled by default.

Issues Fixed in Impala for CDH 5.14.4

For the full list of fixed issues for all CDH components in CDH 5.14.4, see [Issues Fixed in CDH 5.14.x](#) on page 2076. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-4315](#) - Users with the column-level privileges can run the `USE` and `SHOW TABLES` statements.
- [IMPALA-5270](#) - Impala now passes fully resolved sort expressions to prevent crashes with inline view and analytic functions.
- [IMPALA-6500](#) - Fixed an issue where Impala crashed under certain hypervisors that returned out-of-range CPU IDs.
- [IMPALA-6687](#) - The `INSERT` statements with mixed case partition column name no longer fails.
- [IMPALA-6822](#) - A new query option, `SHUFFLE_DISTINCT_EXPRS`, is available for disabling shuffling by distinct expressions.
- [IMPALA-6908](#) - The `ECONNRESET` error code handling was added to the `IsConnReset` function that checks if a given exception is due to a stale connection.
- [IMPALA-6934](#) - Correct results returns from the `EXISTS` subquery containing `ORDER BY`, `LIMIT`, and `OFFSET`
- [IMPALA-6954](#) - Fixed a problem with expression rewrites that caused losses of partitions with CTAS into Kudu.
- [IMPALA-7032](#) - Codegen for `CHAR` type null literals are disabled to prevent a crash.

Issues Fixed in Impala for CDH 5.14.2

For the full list of fixed issues for all CDH components in CDH 5.14.2, see [Issues Fixed in CDH 5.14.x](#) on page 2076. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-3942](#) - Fixed incorrectly escaped string literal in front-end.
- [IMPALA-4664](#) - Fixed an unexpected string conversion in Shell.
- [IMPALA-5139](#) - Updated `mvn-quiet.sh` to print execution content to log file.
- [IMPALA-5269](#) - Fix the issue with final line of query followed by a comment.
- [IMPALA-5936](#) - The `'%'` operator no longer overflows on large decimals.
- [IMPALA-5940](#) - Impala no longer shows stack tracing with `Status::Expected()`.
- [IMPALA-6008](#) - Creating a UDF from a shared library with a `.ll` extension does not crash Impala.
- [IMPALA-6114](#) - Require type equality for `NumericLiteral::localEquals()`.
- [IMPALA-6284](#) - Mark the intermediate decimal avg struct as packed.
- [IMPALA-6295](#) - Fix mix/max handling of `'nan'` and `'inf'`.
- [IMPALA-6307](#) - CTAS statement fails with duplicate column exception.
- [IMPALA-6348](#) - Redact only sensitive fields in runtime profiles.

- [IMPALA-6353](#) - Impala no longer crashes in snappy decompressor.
- [IMPALA-6384](#) - RequestPoolService should honor custom group mapping config.
- [IMPALA-6388](#) - Fix the Union node number of hosts estimation
- [IMPALA-6418](#) - Find a reliable way to detect supported TLS versions.
- [IMPALA-6435](#) - Disabled codegen for CHAR literals.
- [IMPALA-6454](#) - CTAS into Kudu fails with mixed-case partition or primary key column names.
- [IMPALA-6473](#) - An analytic sorts with the same expression in PARTITION BY and ORDER BY does not return an error.
- [IMPALA-6489](#) - Impala now uses the correct template tuple size.
- [IMPALA-6530](#) - Impala now tracks time spent opening HDFS file handles.
- [IMPALA-6538](#) - Fixed read path when Parquet min/max statistics contain NaN.
- [IMPALA-6542](#) - Fixed inconsistent write path of Parquet min/max statistics.
- [IMPALA-6570](#) - Support KUDU on SLES12.
- [IMPALA-6577](#) - Avoid slow Status constructor on expiration thread.

Issues Fixed in Impala for CDH 5.14.0

For the full list of fixed issues for all CDH components in CDH 5.14, see [Issues Fixed in CDH 5.14.x](#) on page 2076. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-467](#) - enable disabled metric verification
- [IMPALA-992](#) - [DOCS] Document impala-shell 'rerun' command
- [IMPALA-1291](#) - Parquet read fails if io buffer size is less than the footer size
- [IMPALA-1575](#) - part 2: yield admission control resources
- [IMPALA-1575](#) - part 2: yield admission control resources
- [IMPALA-1575](#) - Part 1: eagerly release query exec resources
- [IMPALA-1767](#) - Adds predicate to test boolean values true, false, unknown.
- [IMPALA-2107](#) - [DOCS] Document base64*code() functions
- [IMPALA-2181](#) - Add query option levels for display
- [IMPALA-2190](#) - [DOCS] from_timestamp() and to_timestamp()
- [IMPALA-2234](#) - remove workaround from stress test
- [IMPALA-2235](#) - Fix current db when shell auto-reconnects
- [IMPALA-2250](#) - Make multiple COUNT(DISTINCT) message state workarounds
- [IMPALA-2281](#) - Replace FNV with FastHash in exchange nodes
- [IMPALA-2494](#) - Support for byte array encoded decimals in Parquet scanner
- [IMPALA-2636](#) - HS2 GetTables() returns TABLE_TYPE as TABLE for VIEW
- [IMPALA-2758](#) - Remove BufferedTupleStream::GetRows
- [IMPALA-2810](#) - Remove column stats restoration when altering table
- [IMPALA-3200](#) - [DOCS] Document user-facing aspects of new buffer pool
- [IMPALA-3360](#) - Codegen inserting into runtime filters
- [IMPALA-3437](#) - DECIMAL_V2: avoid implicit decimal->double conversion
- [IMPALA-3504](#) - [DOCS] Document utc_timestamp()
- [IMPALA-3516](#) - Avoid writing to /tmp in testing
- [IMPALA-3548](#) - Prune runtime filters based on query options in the FE
- [IMPALA-3613](#) - Avoid topic updates to unregistered subscriber instances
- [IMPALA-3642](#) - Adding backend addresses to error statuses for some scratch failures.
- [IMPALA-3877](#) - support unpatched LLVM
- [IMPALA-3897](#) - Codegen null-aware constant in PHJ::ProcessBuildBatch()
- [IMPALA-3998](#) - deprecate --refresh_after_connect
- [IMPALA-4082](#) - Remove todo item in getRegionsInRange
- [IMPALA-4177](#) - IMPALA-6039: batched bit reading and rle decoding
- [IMPALA-4236](#) - Codegen CopyRows() for select nodes

- [IMPALA-4252](#) - Min-max runtime filters for Kudu
- [IMPALA-4252](#) - Move runtime filters to ScanNode
- [IMPALA-4513](#) - Promote integer types for ABS()
- [IMPALA-4524](#) - Batch ALTER TABLE...ADD PARTITION calls.
- [IMPALA-4591](#) - Bound Kudu client error mem usage
- [IMPALA-4620](#) - Refactor evalcost computation in query analysis
- [IMPALA-4622](#) - [DOCS] New Kudu ALTER TABLE syntax
- [IMPALA-4623](#) - [DOCS] Document file handle caching
- [IMPALA-4669](#) - [KRPC] Add kudu_rpc library to build
- [IMPALA-4670](#) - Introduces RpcMgr class
- [IMPALA-4704](#) - Turns on client connections when local catalog initialized.
- [IMPALA-4736](#) - Add SIGUSR1 behavior to help string for 'minidump_path' flag
- [IMPALA-4786](#) - Clean up how ImpalaServers are created
- [IMPALA-4847](#) - Simplify HdfsTable block metadata loading code
- [IMPALA-4856](#) - Port data stream service to KRPC
- [IMPALA-4856](#) - IMPALA-4872: Include KRPC services in plan fragment's destinations
- [IMPALA-4856](#) - Rename thrift-deps to gen-deps
- [IMPALA-4863](#) - IMPALA-5311: Correctly account the file type and compression codec
- [IMPALA-4918](#) - Support getting column comments via HS2
- [IMPALA-4927](#) - Impala should handle invalid input from Sentry
- [IMPALA-4939](#) - IMPALA-4940: Decimal V2 multiplication
- [IMPALA-4964](#) - Fix Decimal modulo overflow
- [IMPALA-4985](#) - use parquet stats of nested types for dynamic pruning
- [IMPALA-5018](#) - Error on decimal modulo or divide by zero
- [IMPALA-5019](#) - Decimal V2 addition
- [IMPALA-5129](#) - Use Kudu's Kinit code to avoid expensive fork
- [IMPALA-5142](#) - EventSequence displays negative elapsed time.
- [IMPALA-5174](#) - Suppress kudu flags that aren't relevant to Impala
- [IMPALA-5199](#) - prevent hang on empty row batch exchange
- [IMPALA-5210](#) - Count rows and collection items in parquet scanner separately
- [IMPALA-5211](#) - Simplifying nullif conditional.
- [IMPALA-5211](#) - Simplifying ifnull/isnull/nvl where conditional is a literal.
- [IMPALA-5243](#) - Speed up code gen for wide Avro tables.
- [IMPALA-5250](#) - Unify decompressor output_length semantics
- [IMPALA-5307](#) - Part 4: copy out uncompressed text and seq
- [IMPALA-5307](#) - Part 2: copy out strings in uncompressed Avro
- [IMPALA-5307](#) - part 1: don't transfer disk I/O buffers out of parquet
- [IMPALA-5307](#) - Part 3: remove TODO from RCFile
- [IMPALA-5317](#) - add DATE_TRUNC() function
- [IMPALA-5341](#) - Avoid unintended filter out in fe test
- [IMPALA-5376](#) - Implement all TPCDS test cases or alternates for Impala.
- [IMPALA-5383](#) - [DOCS] Document unpartitioned Kudu tables
- [IMPALA-5394](#) - Change ThriftServer() to always use TAcceptQueueServer
- [IMPALA-5416](#) - Fix an impala-shell command recursion bug
- [IMPALA-5417](#) - make I/O buffer queue fixed-size
- [IMPALA-5425](#) - Add test for validating input when setting query options
- [IMPALA-5429](#) - Multi threaded block metadata loading
- [IMPALA-5448](#) - fix invalid number of splits reported in Parquet scan node
- [IMPALA-5501](#) - [DOCS] Clarify 'binaries' w.r.t. installation procedure
- [IMPALA-5525](#) - Extend TestScannersFuzzing to test uncompressed parquet

- [IMPALA-5529](#) - [DOCS] New trunc() signatures
- [IMPALA-5538](#) - Revert "Use explicit catalog versions for deleted objects"
- [IMPALA-5538](#) - Use explicit catalog versions for deleted objects
- [IMPALA-5541](#) - Reject BATCH_SIZE greater than 65536
- [IMPALA-5589](#) - Re-apply:change "set" in impala-shell to show empty string for unset query options
- [IMPALA-5589](#) - Revert "change "set" in impala-shell to show empty string for unset query options"
- [IMPALA-5589](#) - change "set" in impala-shell to show empty string for unset query options
- [IMPALA-5599](#) - Clean up references to TimestampValue in be/src.
- [IMPALA-5599](#) - Fix for mis-use of TimestampValue
- [IMPALA-5624](#) - Replace "ls -l" with opendir() in ProcessStateInfo
- [IMPALA-5638](#) - [DOCS] Add known issue for Impala-Kudu-Sentry issue
- [IMPALA-5653](#) - Remove "unlimited" process mem_limit option
- [IMPALA-5664](#) - Unix time to timestamp conversions may crash Impala
- [IMPALA-5668](#) - Fix cast(X as timestamp) for negative subsecond Decimals
- [IMPALA-5736](#) - Add impala-shell argument to set default query options
- [IMPALA-5789](#) - Add always_false flag in bloom filter
- [IMPALA-5836](#) - Improvements to Eclipse frontend configuration.
- [IMPALA-5844](#) - use a MemPool for expr result allocations
- [IMPALA-5846](#) - Fix output path for kudu libraries
- [IMPALA-5854](#) - Revert "Update external hadoop versions"
- [IMPALA-5854](#) - Update external hadoop versions
- [IMPALA-5860](#) - upgrade to LLVM 3.9.1
- [IMPALA-5870](#) - Improve runtime profile for partial sort
- [IMPALA-5873](#) - Check for existence of sync_file_range()
- [IMPALA-5894](#) - [DOCS] Clarify placement of STRAIGHT_JOIN hint
- [IMPALA-5895](#) - clean up runtime profile lifecycle
- [IMPALA-5902](#) - add ThreadSanitizer build
- [IMPALA-5905](#) - build-all-flag-combinations addendum
- [IMPALA-5905](#) - add script for all-build-options job
- [IMPALA-5908](#) - Allow SET to unset modified query options.
- [IMPALA-5912](#) - fix crash in trunc(..., "WW") in release build
- [IMPALA-5920](#) - addendum - add missing RAT check
- [IMPALA-5920](#) - Remove admission control dependency on YARN RM jar
- [IMPALA-5926](#) - Avoid printing expensive stack when closing a session
- [IMPALA-5927](#) - Fix enable_distcc for zsh
- [IMPALA-5932](#) - Improve transitive closure computation performance in FE
- [IMPALA-5940](#) - Avoid log spew by using Status::Expected()
- [IMPALA-5940](#) - Avoid log spew by using Status::Expected.
- [IMPALA-5941](#) - Fix Metastore schema creation in create-test-configuration.sh
- [IMPALA-5957](#) - print memory address, not memory
- [IMPALA-5965](#) - avoid per-value switch on NeedsConversionInline() in parquet
- [IMPALA-5966](#) - Fix the result file location of PlannerTest
- [IMPALA-5975](#) - Work around broken beeline clients
- [IMPALA-5976](#) - Remove equivalence class computation in FE
- [IMPALA-5986](#) - Correct set-option logic to recognize digits in names.
- [IMPALA-5987](#) - LZ4 Codec silently produces bogus compressed data for large inputs
- [IMPALA-5988](#) - optimise MemPool::TryAllocate()
- [IMPALA-5999](#) - Fix LLVM linkage errors due LibCache sync issues
- [IMPALA-6001](#) - Part 1: Increase log level during catalog update processing
- [IMPALA-6002](#) - Add a LLVM diagnostic handler for LLVM linker errors

- [IMPALA-6004](#) - Fix test_row_filters failure on ASAN
- [IMPALA-6009](#) - Upgrade Guava to 14.0.1
- [IMPALA-6012](#) - workaround - downgrade hive temporarily
- [IMPALA-6016](#) - Fix logging in TableLoadingMgr class
- [IMPALA-6019](#) - Remove dead code parallel-executor*
- [IMPALA-6021](#) - Revert "IMPALA-6009: Upgrade Guava to 14.0.1"
- [IMPALA-6027](#) - Retry downloading toolchain components.
- [IMPALA-6030](#) - Don't start coordinator specific thread pools if a node isn't a coordinator node
- [IMPALA-6040](#) - skip test_multi_compression_types where Hive isn't supported
- [IMPALA-6053](#) - Fix exception when storadelds don't match hosts
- [IMPALA-6054](#) - Parquet dictionary pages should be freed on dictionary construction
- [IMPALA-6055](#) - Fix hdfs encryption test for Hadoop 2.8+
- [IMPALA-6060](#) - Check the return value of JNI exception handling functions
- [IMPALA-6067](#) - Enable S3 access via IAM roles for EC2 VMs
- [IMPALA-6068](#) - Revert "Fix dataload for complextypes_fileformat"
- [IMPALA-6068](#) - Fix dataload for complextypes_fileformat
- [IMPALA-6069](#) - Fix CodegenAnyVal's handling of 'nan'
- [IMPALA-6076](#) - Parquet BIT_PACKED deprecation warning
- [IMPALA-6080](#) - clean up table descriptor handling
- [IMPALA-6081](#) - Fix test_basic_filters runtime profile failure
- [IMPALA-6084](#) - Avoid using of global namespace for llvm
- [IMPALA-6099](#) - Fix crash in CheckForAlwaysFalse()
- [IMPALA-6100](#) - increase test_exchange_delays timeout on slow builds
- [IMPALA-6105](#) - Clarify argument order in single_node_perf_run
- [IMPALA-6106](#) - handle comments before set in test parser
- [IMPALA-6108](#) - IMPALA-6070: Parallel data load (re-instated).
- [IMPALA-6108](#) - Revert "IMPALA-6070: Parallel data load."
- [IMPALA-6109](#) - xfail TestHdfsUnknownErrors::test_hdfs_safe_mode_error_255
- [IMPALA-6118](#) - Fix assertion failure in coordinator bloom filter updating
- [IMPALA-6121](#) - remove I/O mgr request context cache
- [IMPALA-6123](#) - Fix column order of a query test in test_inline_view_limit
- [IMPALA-6124](#) - Fix alter table ddl updates and test
- [IMPALA-6126](#) - ASAN detects heap-use-after-free in thrift-server-test
- [IMPALA-6127](#) - Fix timeout in TestRuntimeFilter.test_wait_time
- [IMPALA-6128](#) - Spill-to-disk Encryption(AES-CFB + SHA256) is slow
- [IMPALA-6134](#) - Update code base to use impala::ConditionVariable
- [IMPALA-6136](#) - Part 1: Query duration should not be normally negative.
- [IMPALA-6137](#) - fix text scanner split delim mem mgmt
- [IMPALA-6144](#) - PublishFilter() continues to run after query failure/cancellation
- [IMPALA-6148](#) - Specifying thirdparty deps as URLs
- [IMPALA-6151](#) - add query-level fragment/backend counters
- [IMPALA-6155](#) - Allow tests to pass when ORDER BY does not cover the query.
- [IMPALA-6164](#) - Fix stale query profile in TestAlwaysFalseFilter
- [IMPALA-6171](#) - Revert "IMPALA-1575: part 2: yield admission control resources"
- [IMPALA-6173](#) - Fix SHOW CREATE TABLE for unpartitioned Kudu tables
- [IMPALA-6183](#) - Fix Decimal to Double conversion
- [IMPALA-6184](#) - Clean up after ScalarExprEvaluator::Clone() fails
- [IMPALA-6187](#) - Fix missing conjuncts evaluation with empty projection
- [IMPALA-6188](#) - make test_top_n_reclaim less flaky
- [IMPALA-6198](#) - marks a test as debug-only

- [IMPALA-6201](#) - Fix test_basic_filters on ASAN
- [IMPALA-6206](#) - Fix data load failure with -notests
- [IMPALA-6217](#) - fix DCHECK in Parquet fuzz test
- [IMPALA-6220](#) - Revert "IMPALA-6128: Spill-to-disk Encryption(AES-CFB + SHA256) is slow"
- [IMPALA-6225](#) - Part 1: Query profile date-time strings should have ns precision.
- [IMPALA-6232](#) - Disable file handle cache by default
- [IMPALA-6241](#) - timeout in admission control test under ASAN
- [IMPALA-6256](#) - Incorrect principal will be used for internal connections if FLAGS_be_principal is set
- [IMPALA-6262](#) - Always initialize runtime profile for DataSink
- [IMPALA-6280](#) - Materialize TupleIsNullPredicate for insert sorts
- [IMPALA-6281](#) - Fix use-after-free in InitAuth()
- [IMPALA-6285](#) - Don't print stack trace on RPC errors.
- [IMPALA-6286](#) - Remove invalid runtime filter targets.
- [IMPALA-6291](#) - disable AVX512 codegen in LLVM
- [IMPALA-6292](#) - addendum: Fix failing test
- [IMPALA-6292](#) - Fix incorrect DCHECK in decimal subtraction
- [IMPALA-6298](#) - Skip test_profile_fragment_instances on local filesystem
- [IMPALA-6308](#) - Fix bad Status() usage in data-stream-sender.cc
- [IMPALA-5307](#) - ScannerContext API change
- [IMPALA-5493](#) - Add Protobuf to CMakeLists.txt

Issues Fixed in Impala for CDH 5.13.3

For the full list of fixed issues for all CDH components in CDH 5.13.3, see [Issues Fixed in CDH 5.13.x](#) on page 2106. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-6353](#) - Fixed the crash in snappy decompressor.
- [IMPALA-6473](#) - Error in analytic sort with same expr in 'partition by' and 'order by'.
- [IMPALA-6500](#) - Impala crashes randomly on different queries with GROUP BY.

Issues Fixed in Impala for CDH 5.13.2

For the full list of fixed issues for all CDH components in CDH 5.13.2, see [Issues Fixed in CDH 5.13.x](#) on page 2106. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-2810](#) - Remove column stats restoration when altering table.
- [IMPALA-4664](#) - Unexpected string conversion in Shell.
- [IMPALA-5167](#) - Revert "Reduce the number of Kudu clients created (FE)".
- [IMPALA-5664](#) - Unix time to timestamp conversions may crash Impala
- [IMPALA-5926](#) - Avoid printing expensive stack when closing a session.
- [IMPALA-5936](#) - Operator '%' overflows on large decimals.
- [IMPALA-5940](#) - Avoid log spew by using Status::Expected.
- [IMPALA-5951](#) - Remove flaky test_catalogd_timeout.
- [IMPALA-5987](#) - LZ4 Codec silently produces bogus compressed data for large inputs.
- [IMPALA-6004](#) - Fix the test_row_filters failure on ASAN.
- [IMPALA-6053](#) - Fix the exception when storadelds don't match hosts.
- [IMPALA-6060](#) - Check the return value of JNI exception handling functions.
- [IMPALA-6067](#) - Enable S3 access via IAM roles for EC2 VMs.
- [IMPALA-6069](#) - Fix CodegenAnyVal's handling of 'nan'.
- [IMPALA-6100](#) - Increase test_exchange_delays timeout on slow builds.
- [IMPALA-6106](#) - Handle comments before set in test parser.
- [IMPALA-6114](#) - Require type equality for NumericLiteral::localEquals().
- [IMPALA-6137](#) - Fix text scanner split delim mem mgmt.
- [IMPALA-6183](#) - Fix Decimal to Double conversion.

- [IMPALA-6184](#) - Clean up after ScalarExprEvaluator::Clone() fails.
- [IMPALA-6280](#) - Materialize TupleIsNullPredicate for insert sorts.
- [IMPALA-6284](#) - Mark the intermediate decimal avg struct as packed.
- [IMPALA-6286](#) - Remove invalid runtime filter targets.
- [IMPALA-6291](#) - Disable AVX512 codegen in LLVM.
- [IMPALA-6295](#) - Fix mix/max handling of 'nan' and 'inf'.
- [IMPALA-6307](#) - CTAS statement fails with duplicate column exception.
- [IMPALA-6308](#) - Fix bad Status() usage in data-stream-sender.cc.
- [IMPALA-6348](#) - Redact only sensitive fields in runtime profiles.
- [IMPALA-6364](#) - Bypass file handle cache for ineligible files.
- [IMPALA-6381](#) - Increase test_exchange_delays timeout for isilon.
- [IMPALA-6384](#) - RequestPoolService should honor custom group mapping config.
- [IMPALA-6388](#) - Fix the Union node number of hosts estimation.
- [IMPALA-6418](#) - Find a reliable way to detect supported TLS versions.
- [IMPALA-6435](#) - Disable codegen for CHAR literals.

Issues Fixed in Impala for CDH 5.13.1

For the full list of fixed issues for all CDH components in CDH 5.13.1, see [Issues Fixed in CDH 5.13.x](#) on page 2106. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-4682](#) - Remove Preconditions check from analyzeAggregation().
- [IMPALA-4951](#) - Fix database visibility for user with only column privilege
- [IMPALA-5597](#) - Try casting targetExpr when building runtime filter plan
- [IMPALA-5856](#) - Fix outer join predicate assignment.
- [IMPALA-5923](#) - Print binary ID as hex in ChildQuery::Cancel()
- [IMPALA-5954](#) - Set DO_NOT_UPDATE_STATS in alterTable() RPCs to HMS.
- [IMPALA-5955](#) - Use totalSize tblproperty instead of rawDataSize.
- [IMPALA-5983](#) - Fix crash in to/from_utc_timestamp

Issues Fixed in Impala for CDH 5.13.0

For the full list of fixed issues for all CDH components in CDH 5.13, see [Issues Fixed in CDH 5.13.x](#) on page 2106. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-992](#) - Rerun past queries from history in shell
- [IMPALA-1470](#) - Fix error message with catalog down
- [IMPALA-1478](#) - Improve error message when subquery is used in the ON clause
- [IMPALA-1882](#) - Remove ORDER BY restriction from first_value()/last_value()
- [IMPALA-1891](#) - Statestore won't send deletions in initial non-delta topic
- [IMPALA-2373](#) - Extrapolate row counts for HDFS tables.
- [IMPALA-2525](#) - Treat parquet ENUMs as STRINGs when creating impala tables.
- [IMPALA-2615](#) - support [[nodiscard]] on Status
- [IMPALA-3040](#) - addendum: use specific_build_type_timeout for slow builds
- [IMPALA-3208](#) - max_row_size option
- [IMPALA-3504](#) - UDF for current timestamp in UTC
- [IMPALA-3894](#) - Change the behavior parsing date "YY"
- [IMPALA-3905](#) - HdfsScanner::GetNext() for Avro, RC, and Seq scans.
- [IMPALA-3931](#) - arbitrary fixed-size uda intermediate types
- [IMPALA-3937](#) - Deprecate --be_service_threads
- [IMPALA-4039](#) - Increase width of Operator column in runtime profile
- [IMPALA-4086](#) - Add benchmark for simple scheduler
- [IMPALA-4164](#) - Avoid overly aggressive inlining in LLVM IR
- [IMPALA-4192](#) - Disentangle Expr and ExprContext

- [IMPALA-4276](#) - Profile displays non-default query options set by planner
- [IMPALA-4418](#) - Fixes extra blank lines in query result
- [IMPALA-4622](#) - Add ALTER COLUMN statement.
- [IMPALA-4623](#) - Enable file handle cache
- [IMPALA-4669](#) - Add Kudu's RPC, util, and security libraries
- [IMPALA-4674](#) - Port spilling ExecNodes to new buffer pool
- [IMPALA-4687](#) - Get Impala working against HBase 2.0
- [IMPALA-4703](#) - reservation denial debug action
- [IMPALA-4737](#) - Prevent SIGUSR1 from killing daemons when minidumps are disabled
- [IMPALA-4794](#) - Grouping distinct agg plan robust to data skew
- [IMPALA-4795](#) - Allow fetching function obj from catalog using signature
- [IMPALA-4826](#) - Fix error during a scan on repeated root schema in Parquet.
- [IMPALA-4833](#) - Use scheduling information to make per-node memory reservation tight
- [IMPALA-4861](#) - READ_WRITE warning on CREATE TABLE LIKE PARQUET.
- [IMPALA-4862](#) - make resource profile consistent with backend behaviour
- [IMPALA-4866](#) - Hash join node does not apply limits correctly
- [IMPALA-4905](#) - Don't send empty insert status to coordinator
- [IMPALA-4925](#) - Cancel finstance if query has finished
- [IMPALA-5016](#) - Simplify COALESCE() in SimplifyConditionalsRule.
- [IMPALA-5031](#) - UBSAN clean and method for testing UBSAN cleanliness
- [IMPALA-5036](#) - Parquet count star optimization
- [IMPALA-5061](#) - Populate null_count in parquet::statistics
- [IMPALA-5085](#) - large rows in BufferedTupleStreamV2
- [IMPALA-5104](#) - Admit queries with mem equal to proc mem_limit
- [IMPALA-5108](#) - idle_session_timeout kicks in later than expected
- [IMPALA-5109](#) - Increase range of backend latency histogram
- [IMPALA-5116](#) - Remove deprecated hash_* types in gutil
- [IMPALA-5158](#) - Account for difference between process memory consumption and memory used by queries
- [IMPALA-5160](#) - adjust spill buffer size based on planner estimates
- [IMPALA-5167](#) - Reduce the number of Kudu clients created
- [IMPALA-5240](#) - Allow config of number of disk I/O threads per disk type
- [IMPALA-5275](#) - Avoid printing status stack trace on hot paths
- [IMPALA-5280](#) - Coalesce chains of OR conditions to an IN predicate
- [IMPALA-5286](#) - Query with Kudu col name w/ different casing from 'order by' fails
- [IMPALA-5316](#) - Adds last_day() function
- [IMPALA-5327](#) - Handle return of JNI GetStringUTFChar
- [IMPALA-5336](#) - Fix partition pruning when column is cast
- [IMPALA-5347](#) - reduce codegen overhead of timestamp trunc()
- [IMPALA-5350](#) - Tidy up thread groups for first exec threads
- [IMPALA-5352](#) - Age out unused file handles from the cache
- [IMPALA-5354](#) - INSERT hints for Kudu tables
- [IMPALA-5376](#) - Loads all TPC-DS tables
- [IMPALA-5386](#) - Fix ReopenCachedHdfsFileHandle failure case
- [IMPALA-5389](#) - simplify BufferDescriptor lifetime
- [IMPALA-5407](#) - Fix crash in HdfsSequenceTableWriter
- [IMPALA-5412](#) - Fix scan result with partitions on same file
- [IMPALA-5420](#) - Skip ACL fetch if the acl bit is not set.
- [IMPALA-5427](#) - Fix race between CRS::UpdateQueryStatus() and beeswax RPCs
- [IMPALA-5431](#) - Remove redundant path exists checks during table load
- [IMPALA-5433](#) - Mark single-argument Status c'tors as explicit

- [IMPALA-5446](#) - dropped Sorter::Reset() status
- [IMPALA-5477](#) - Fix minidump-2-core tool
- [IMPALA-5479](#) - Propagate the argument type in RawValue::Compare()
- [IMPALA-5480](#) - Improve missing filters message
- [IMPALA-5481](#) - Clarify RowDescriptor ownership
- [IMPALA-5482](#) - fix git checkout when workloads are modified
- [IMPALA-5483](#) - Automatically disable codegen for small queries
- [IMPALA-5484](#) - Fix LICENSE issues discovered by IPMC in 2.9 vote
- [IMPALA-5488](#) - Fix handling of exclusive HDFS file handles
- [IMPALA-5489](#) - Improve Sentry authorization for Kudu tables
- [IMPALA-5492](#) - Fix incorrect newline character in the LDAP message
- [IMPALA-5494](#) - Fixes the selectivity of NOT IN predicates
- [IMPALA-5495](#) - Improve error message if no impalad role is configured
- [IMPALA-5497](#) - spilling hash joins that output build rows hit OOM
- [IMPALA-5498](#) - Support for partial sorts in Kudu INSERTS
- [IMPALA-5499](#) - avoid ephemeral port conflicts
- [IMPALA-5500](#) - Reduce catalog update topic size
- [IMPALA-5504](#) - Fix TupleIsNullPredicate evaluation.
- [IMPALA-5506](#) - Add stdin description to help information of query_file option
- [IMPALA-5507](#) - Add clear description to help information of KEYVAL option
- [IMPALA-5511](#) - Add process start time to debug web page
- [IMPALA-5513](#) - Fix display message exception when using invalid KEYVAL
- [IMPALA-5514](#) - Throw an error when --ldap_password_cmd is used without LDAP auth
- [IMPALA-5517](#) - Allow IMPALA_LOGS_DIR to be overridden
- [IMPALA-5520](#) - TopN node periodically reclaims old allocations
- [IMPALA-5524](#) - Fixes NPE during planning with DISABLE_UNSAFE_SPILLS=1
- [IMPALA-5529](#) - Add additional function signatures for TRUNC()
- [IMPALA-5531](#) - Fix correctness issue in correlated aggregate subqueries
- [IMPALA-5532](#) - Stack-allocate compressors in RowBatch (de)serialization
- [IMPALA-5536](#) - Fix TCLIService thrift compilation on Hive 2
- [IMPALA-5539](#) - Fix Kudu timestamp with -use_local_tz_for_unix_ts
- [IMPALA-5540](#) - Revert Sentry version back to 5.13
- [IMPALA-5546](#) - Allow creating unpartitioned Kudu tables
- [IMPALA-5547](#) - Rework FK/PK join detection.
- [IMPALA-5548](#) - Fix some minor nits with HDFS parquet column readers
- [IMPALA-5549](#) - Remove deprecated fields from CatalogService API
- [IMPALA-5551](#) - Fix AggregationNode::Close() when Prepare() fails
- [IMPALA-5554](#) - sorter DCHECK on null column
- [IMPALA-5560](#) - always store CHAR(N) inline in tuple
- [IMPALA-5567](#) - race in fragment instance teardown
- [IMPALA-5570](#) - fix spilling null-aware anti join
- [IMPALA-5572](#) - Timestamp codegen for text scanner
- [IMPALA-5573](#) - Add decimal codegen in text scanner
- [IMPALA-5579](#) - Fix IndexOutOfBoundsException in GetTables metadata request
- [IMPALA-5580](#) - fix Java UDFs that return NULL strings
- [IMPALA-5582](#) - Store sentry privileges in lower case
- [IMPALA-5586](#) - Null-aware anti-join can take a long time to cancel
- [IMPALA-5588](#) - Reduce the frequency of fault injection
- [IMPALA-5591](#) - set should handle negative values
- [IMPALA-5594](#) - don't import shaded classes

- [IMPALA-5595](#) - Only set KuduScanner timestamp feature flag if necessary
- [IMPALA-5598](#) - Fix excessive dumping in MemLimitExceeded
- [IMPALA-5602](#) - Fix query optimization for kudu and datasource tables
- [IMPALA-5611](#) - KuduPartitionExpr holds onto memory unnecessarily
- [IMPALA-5612](#) - join inversion should factor in parallelism
- [IMPALA-5615](#) - Fix compute incremental stats for general partition exprs
- [IMPALA-5616](#) - Add --enable_minidumps startup flag
- [IMPALA-5617](#) - Include full workload name in tpch_nested query filenames
- [IMPALA-5618](#) - buffered-tuple-stream-v2 fixes
- [IMPALA-5623](#) - Fix lag() on STRING cols to release UDF mem
- [IMPALA-5627](#) - fix dropped statuses in HDFS writers
- [IMPALA-5629](#) - avoid expensive list::size() call
- [IMPALA-5630](#) - Add Kudu client version as a common metric
- [IMPALA-5636](#) - Change the metadata in parquet
- [IMPALA-5638](#) - Fix Kudu table set tblproperties inconsistencies
- [IMPALA-5641](#) - mem-estimate should never be less than mem-reservation
- [IMPALA-5643](#) - Add total number of threads created per group to /threadz
- [IMPALA-5644](#) - Fail queries early when their minimum reservation is too high to execute within the given mem_limit
- [IMPALA-5648](#) - fix count(*) mem estimate regression
- [IMPALA-5650](#) - Make sum_init_zero a SUM function
- [IMPALA-5652](#) - deprecate unlimited process mem_limit
- [IMPALA-5657](#) - FunctionCallExpr.toSql() and clone() ignore 'IGNORE NULLS' case
- [IMPALA-5658](#) - addtl. process/system-wide memory metrics
- [IMPALA-5659](#) - Begin standardizing treatment of thirdparty libraries
- [IMPALA-5659](#) - glog/gflags should be dynamically linked
- [IMPALA-5661](#) - buffer pool limit
- [IMPALA-5666](#) - ASAN poisoning for MemPool and BufferPool
- [IMPALA-5670](#) - Misc. tidying of ExecEnv
- [IMPALA-5676](#) - avoid expensive consistency checks in BTSv2
- [IMPALA-5677](#) - limit clean page memory consumption
- [IMPALA-5679](#) - Fix Parquet count(*) with group by string
- [IMPALA-5681](#) - release reservation from blocking operators
- [IMPALA-5686](#) - Update a mini cluster Sentry property
- [IMPALA-5689](#) - Avoid inverting non-equi left joins
- [IMPALA-5691](#) - recalibrate mem limit for Q18
- [IMPALA-5696](#) - Enable cipher configuration when using TLS / Thrift
- [IMPALA-5709](#) - Remove mini-impala-cluster
- [IMPALA-5713](#) - always reserve memory for preaggs
- [IMPALA-5714](#) - Add OpenSSL to bootstrap_toolchain.py
- [IMPALA-5715](#) - (mitigation only) defer destruction of MemTrackers
- [IMPALA-5716](#) - Don't delete cmake_modules/* when enabling distcc
- [IMPALA-5722](#) - Fix string to decimal cast
- [IMPALA-5725](#) - coalesce() with outer join incorrectly rewritten
- [IMPALA-5739](#) - Correctly handle sles12 SP2
- [IMPALA-5742](#) - De-allocate buffer in parquet-reader on exit
- [IMPALA-5743](#) - Support TLS version configuration for Thrift servers
- [IMPALA-5744](#) - Add 'use_krpc' flag and create DataStream interface
- [IMPALA-5745](#) - Bump Breakpad version
- [IMPALA-5749](#) - coordinator race hits DCHECK 'num_remaining_backends_ > 0'
- [IMPALA-5750](#) - Catch exceptions from boost thread creation

- [IMPALA-5756](#) - start memory maintenance thread after metric creation
- [IMPALA-5757](#) - Make tbl property toSql deterministic
- [IMPALA-5769](#) - Add periodic minidump cleanup
- [IMPALA-5773](#) - Correctly account for memory used in data stream receiver queue
- [IMPALA-5774](#) - Prevent FindInSet() from reading off end of string
- [IMPALA-5775](#) - Allow shell to support TLSv1, v1.1 and v1.2
- [IMPALA-5776](#) - HdfsTextScanner::WritePartialTuple() writes the varlen data to an incorrect memory pool
- [IMPALA-5778](#) - Clarify logging around and usage of read_size startup option
- [IMPALA-5784](#) - Separate planner and user set query options in profile
- [IMPALA-5787](#) - Dropped status in KuduTableSink::Send()
- [IMPALA-5788](#) - Fix agg node crash when grouping by nondeterministic exprs
- [IMPALA-5791](#) - Make bootstrap_development.sh survive apt-get failure
- [IMPALA-5796](#) - CTAS for Kudu fails with expr rewrite
- [IMPALA-5798](#) - ASAN use-after-poison in Parquet decoder
- [IMPALA-5799](#) - Kudu DML can crash if schema has changed
- [IMPALA-5800](#) - Configure Squeasel's cipher suite and TLS version
- [IMPALA-5811](#) - Add 'backends' tab to query details pages
- [IMPALA-5812](#) - Fix NPE when joining on empty const select
- [IMPALA-5815](#) - right outer join returns invalid memory
- [IMPALA-5819](#) - DCHECK in HdfsTextScanner::Close()
- [IMPALA-5823](#) - fix SET_DENY_RESERVATION_PROBABILITY
- [IMPALA-5825](#) - Catch exceptions thrown by TSSLocketFactory c'tor
- [IMPALA-5838](#) - Improve errors on AC buffer mem rejection
- [IMPALA-5840](#) - Don't write page-level statistics in Parquet files.
- [IMPALA-5849](#) - Remove compile-time checks for OpenSSL > 1.0.0
- [IMPALA-5850](#) - Cast sender partition exprs under unions.
- [IMPALA-5852](#) - improve MINIMUM_RESERVATION_UNAVAILABLE error
- [IMPALA-5853](#) - fix GetResultSetMetadata() error message for invalid query id
- [IMPALA-5855](#) - Preaggregation crashes - unable to initialize hash table
- [IMPALA-5857](#) - avoid invalid free of hedged read metrics
- [IMPALA-5867](#) - Fix bugs parsing 2-digit year
- [IMPALA-5871](#) - KuduPartitionExpr incorrectly handles its child types
- [IMPALA-5885](#) - free runtime filter allocations in Parquet
- [IMPALA-5888](#) - free other local allocations in Parquet
- [IMPALA-5890](#) - Abort queries if scanner hits IO errors
- [IMPALA-5891](#) - fix PeriodicCounterUpdater initialization
- [IMPALA-5892](#) - Allow reporting status independent of fragment instance
- [IMPALA-2615](#) - annotate Status with [[nodiscard]]
- [IMPALA-5412](#) - Fix scan result with partitions on same file

Issues Fixed in Impala for CDH 5.12.2

For the full list of fixed issues for all CDH components in CDH 5.12, see [Issues Fixed in CDH 5.12.x](#) on page 2134. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-4682](#) - Remove Preconditions check from analyzeAggregation().
- [IMPALA-4737](#) - Prevent SIGUSR1 from killing daemons when minidumps are disabled
- [IMPALA-4826](#) - Fix error during a scan on repeated root schema in Parquet.
- [IMPALA-4951](#) - Fix database visibility for user with only column privilege
- [IMPALA-5327](#) - Handle return of JNI GetStringUTFChar
- [IMPALA-5336](#) - Fix partition pruning when column is cast
- [IMPALA-5446](#) - dropped Sorter::Reset() status

- [IMPALA-5495](#) - Improve error message if no impalad role is configured
- [IMPALA-5504](#) - Fix TupleIsNullPredicate evaluation.
- [IMPALA-5531](#) - Fix correctness issue in correlated aggregate subqueries
- [IMPALA-5595](#) - Only set KuduScanner timestamp feature flag if necessary
- [IMPALA-5597](#) - Try casting targetExpr when building runtime filter plan
- [IMPALA-5598](#) - Fix excessive dumping in MemLimitExceeded
- [IMPALA-5689](#) - Avoid inverting non-equi left joins
- [IMPALA-5750](#) - Catch exceptions from boost thread creation
- [IMPALA-5776](#) - Write partial tuple to the correct mempool
- [IMPALA-5784](#) - Separate planner and user set query options in profile
- [IMPALA-5796](#) - CTAS for Kudu fails with expr rewrite
- [IMPALA-5815](#) - right outer join returns invalid memory
- [IMPALA-5819](#) - DCHECK in HdfsTextScanner::Close()
- [IMPALA-5820](#) - Fix string format() syntax in test_scanners_fuzz.py
- [IMPALA-5840](#) - Don't write page-level statistics in Parquet files.
- [IMPALA-5850](#) - Cast sender partition exprs under unions.
- [IMPALA-5856](#) - Fix outer join predicate assignment.
- [IMPALA-5871](#) - KuduPartitionExpr incorrectly handles its child types
- [IMPALA-5885](#) - free runtime filter allocations in Parquet
- [IMPALA-5892](#) - Allow reporting status independent of fragment instance
- [IMPALA-5923](#) - Print binary ID as hex in ChildQuery::Cancel()
- [IMPALA-5954](#) - Set DO_NOT_UPDATE_STATS in alterTable() RPCs to HMS.
- [IMPALA-5983](#) - Fix crash in to/from_utc_timestamp
- [IMPALA-5994](#) - Lower case struct-field names
- [IMPALA-6060](#) - Check the return value of JNI exception handling functions

Issues Fixed in Impala for CDH 5.12.1

For the full list of fixed issues for all CDH components in CDH 5.12, see [Issues Fixed in CDH 5.12.x](#) on page 2134. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-4276](#) - Profile displays non-default query options set by planner
- [IMPALA-4866](#) - Hash join node does not apply limits correctly
- [IMPALA-5354](#) - INSERT hints for Kudu tables
- [IMPALA-5427](#) - Fix race between CRS::UpdateQueryStatus() and beeswax RPCs
- [IMPALA-5500](#) - Reduce catalog update topic size
- [IMPALA-5524](#) - Fixes NPE during planning with DISABLE_UNSAFE_SPILLS=1
- [IMPALA-5539](#) - Fix Kudu timestamp with -use_local_tz_for_unix_ts
- [IMPALA-5554](#) - sorter DCHECK on null column
- [IMPALA-5567](#) - race in fragment instance teardown
- [IMPALA-5579](#) - Fix IndexOutOfBoundsException in GetTables metadata request
- [IMPALA-5580](#) - fix Java UDFs that return NULL strings
- [IMPALA-5582](#) - Store sentry privileges in lower case
- [IMPALA-5586](#) - Null-aware anti-join can take a long time to cancel
- [IMPALA-5588](#) - Reduce the frequency of fault injection
- [IMPALA-5611](#) - KuduPartitionExpr holds onto memory unnecessarily
- [IMPALA-5615](#) - Fix compute incremental stats for general partition exprs
- [IMPALA-5616](#) - Add --enable_minidumps startup flag
- [IMPALA-5623](#) - Fix lag() on STRING cols to release UDF mem
- [IMPALA-5627](#) - fix dropped statuses in HDFS writers
- [IMPALA-5638](#) - Fix Kudu table set tblproperties inconsistencies
- [IMPALA-5657](#) - Fix a couple of bugs with FunctionCallExpr and IGNORE NULLS

- [IMPALA-5686](#) - Update a mini cluster Sentry property
- [IMPALA-5691](#) - recalibrate mem limit for Q18

Issues Fixed in Impala for CDH 5.12.0

For the full list of fixed issues for all CDH components in CDH 5.12, see [Issues Fixed in CDH 5.12.x](#) on page 2134. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-278](#) - Mention DIV arithmetic operator.
- [IMPALA-1427](#) - Improvements to "Unknown disk-ID" warning
- [IMPALA-1670](#), [IMPALA-4141](#) - Support multiple partitions in ALTER TABLE ADD PARTITION
- [IMPALA-1861](#) - Simplify conditionals with constant conditions
- [IMPALA-1972](#), [IMPALA-3882](#) - Fix client_request_state_map_lock_contention
- [IMPALA-2020](#) - Inline big number strings
- [IMPALA-2079](#) - Don't fail when write to scratch dir results in error.
- [IMPALA-2328](#) - Read support for min/max Parquet statistics
- [IMPALA-2518](#) - DROP DATABASE CASCADE removes cache directives of tables
- [IMPALA-2522](#) - Improve the reliability and effectiveness of ETL.
- [IMPALA-2550](#) - Switch to per-query exec rpc
- [IMPALA-2615](#) - warn if Status is ignored
- [IMPALA-3079](#) - Fix sequence file writer
- [IMPALA-3200](#) - Implement suballocator for splitting buffers
- [IMPALA-3202](#) - implement spill-to-disk in new buffer pool
- [IMPALA-3203](#) - Implement scalable buffer recycling in buffer pool.
- [IMPALA-3524](#) - Don't process spilled partitions with 0 probe rows
- [IMPALA-3586](#) - Implement union passthrough
- [IMPALA-3654](#) - Parquet stats filtering for IN predicate
- [IMPALA-3671](#) - Add SCRATCH_LIMIT query option.
- [IMPALA-3742](#) - Partitions and sort INSERTs for Kudu tables
- [IMPALA-3748](#) - minimum buffer requirements in planner
- [IMPALA-3748](#) - add query-wide resource acquisition step
- [IMPALA-3785](#) - Record query handle for invalid handle
- [IMPALA-3794](#) - Workaround for Breakpad ID conflicts
- [IMPALA-3905](#) - Implements HdfsScanner::GetNext() for text scans.
- [IMPALA-3909](#) - Populate min/max statistics in Parquet writer
- [IMPALA-3989](#) - Display skew warning for poorly formatted Parquet files
- [IMPALA-4008](#) - Don't bake fields into generated IR functions of OldHashTable
- [IMPALA-4014](#) - Introduce query-wide execution state.
- [IMPALA-4029](#) - Reduce memory requirements for storing file metadata
- [IMPALA-4036](#) - invalid SQL generated for partitioned table with comment
- [IMPALA-4041](#) - Limit catalog and admission control updates to coordinators
- [IMPALA-4163](#) - Add sortby() query hint
- [IMPALA-4166](#) - Add SORT BY sql clause
- [IMPALA-4341](#) - Add metadata load to planner timeline
- [IMPALA-4351](#), [IMPALA-4353](#) - query generator random profile options for INSERT.
- [IMPALA-4355](#) - random query generator: modify statement execution flow to support DML
- [IMPALA-4359](#) - qgen: add UPSERT support
- [IMPALA-4390](#) - Separate ADD and DROP PARTITION syntax
- [IMPALA-4431](#) - Add audit event log control mechanism to prevent disk overflow
- [IMPALA-4477](#) - Update Kudu to latest commit on master
- [IMPALA-4499](#) - Table name missing from exec summary
- [IMPALA-4546](#) - Fix Moscow timezone conversion after 2014

- [IMPALA-4548](#) - BlockingJoinNode should wait for async build thread
- [IMPALA-4549](#) - consistently treat 9999 as upper bound for timestamp year
- [IMPALA-4588](#) - Enable starting the minicluster when offline
- [IMPALA-4611](#) - Checking perms on S3 files is a very expensive no-op
- [IMPALA-4616](#) - Add missing Kudu column options
- [IMPALA-4617](#) - remove IsConstant() analysis from backend
- [IMPALA-4624](#) - Implement Parquet dictionary filtering
- [IMPALA-4631](#) - avoid DCHECK in PlanFragementExecutor::Close().
- [IMPALA-4640](#) - Fix number of rows displayed by parquet-reader tool
- [IMPALA-4647](#) - fix full data load with ninja
- [IMPALA-4648](#) - remove build_thirdparty.sh
- [IMPALA-4649](#) - add a mechanism to pass flags into make
- [IMPALA-4650](#) - Add Protobuf 2.6.1 to toolchain and as a build dependency
- [IMPALA-4651](#) - Add LibEv to build
- [IMPALA-4652](#) - Add crcutil to build
- [IMPALA-4653](#) - fix sticky config variable problem
- [IMPALA-4662](#) - Fix NULL literal handling in Kudu IN list predicates
- [IMPALA-4673](#) - Use --local_library_dir for tzdb startup scratch space
- [IMPALA-4674](#) - port BufferedTupleStream to BufferPool
- [IMPALA-4676](#) - remove vestigial references to getBlockStorageLocations() API
- [IMPALA-4678](#) - move query MemTracker into QueryState
- [IMPALA-4684](#) - check-hbase-nodes.py: Build failing on RHEL7 when trying to start HBase.
- [IMPALA-4686](#) - Fix schema output for INT96 columns in parquet-reader tool
- [IMPALA-4689](#) - Fix computation of last active time
- [IMPALA-4701](#) - make distcc work reliably with clang
- [IMPALA-4707](#) - fix use-after-free in QueryExecMgr
- [IMPALA-4710](#) - There is an error in control audit log file size number
- [IMPALA-4711](#) - clarify is_null semantics in udf.h
- [IMPALA-4716](#) - Expr rewrite causes IllegalStateException
- [IMPALA-4731](#), [IMPALA-397](#), [IMPALA-4728](#) - Sorter crash Impalad instance.
- [IMPALA-4734](#) - Set parquet::RowGroup::sorting_columns
- [IMPALA-4738](#) - STDDEV_SAMP should return NULL for single record input
- [IMPALA-4740](#) - Add option to use hdfsPread() for HDFS hedged reads
- [IMPALA-4747](#) - macros should only evaluate their arguments once
- [IMPALA-4748](#) - crash in TmpFileMgr when hitting process mem limit
- [IMPALA-4752](#) - make ObjectPool more efficient
- [IMPALA-4757](#) - addendum: avoid double underscore in name
- [IMPALA-4758](#) - Upgrade gutil to recent Kudu version.
- [IMPALA-4762](#) - RECOVER PARTITIONS should batch partition updates
- [IMPALA-4764](#) - Add Hedged read metrics
- [IMPALA-4788](#) - Use HashSet in RECOVER PARTITIONS duplicate checks
- [IMPALA-4789](#) - Fix slow metadata loading due to inconsistent paths.
- [IMPALA-4792](#) - Fix number of distinct values for a CASE with constant outputs
- [IMPALA-4801](#) - fix heap use after free for MemTracker
- [IMPALA-4815](#) -, [IMPALA-4817](#), [IMPALA-4819](#): Write and Read Parquet Statistics for remaining types
- [IMPALA-4820](#) - avoid writing unencrypted data during write cancellation
- [IMPALA-4828](#) - Add "Known Issues" item for
- [IMPALA-4831](#) - enforce BufferPool reservation invariants
- [IMPALA-4839](#) - Remove implicit 'localhost' for KUDU_MASTER_HOSTS
- [IMPALA-4840](#) - Fix REFRESH performance regression.

- [IMPALA-4846](#) - Upgrade Snappy to 1.1.4
- [IMPALA-4849](#) - IllegalStateException from rewritten CASE expr
- [IMPALA-4858](#) - add more info to MemLimitExceeded errors
- [IMPALA-4859](#) - Push down IS NULL / IS NOT NULL to Kudu
- [IMPALA-4877](#) - fix precedence of unary -/+
- [IMPALA-4880](#) - Clarify synchronization policy for 'done_' in KuduScanNode
- [IMPALA-4883](#) - Union Codegen
- [IMPALA-4884](#) - Add JVM heap and non-heap usage in metrics and UI
- [IMPALA-4885](#) - Expose Jvm thread info in web UI
- [IMPALA-4890](#) - /5143: Coordinator race involving TearDown()
- [IMPALA-4892](#) - Session ID included in error message
- [IMPALA-4893](#) - Efficiently update the rows read counter for sequence file
- [IMPALA-4897](#) - AnalysisException: specified cache pool does not exist
- [IMPALA-4923](#) - reduce memory transfer for selective scans
- [IMPALA-4926](#) - Upgrade LZ4 to 1.7.5
- [IMPALA-4937](#) - Remove unused kudu scanner keep alive variable
- [IMPALA-4943](#) - Speed up block md loading for add/recover partition calls.
- [IMPALA-4946](#) - fix hang in BufferPool
- [IMPALA-4955](#) - Fix integer overflow in hdfs table size accounting
- [IMPALA-4959](#) - Avoid picking up the system's boost cmake module
- [IMPALA-4962](#) - Fix SHOW COLUMN STATS for HS2
- [IMPALA-4965](#) - Authorize access to runtime profile and exec summary
- [IMPALA-4966](#) - Add flatbuffers to build
- [IMPALA-4983](#) - Set toolchain version to include LZ4 build flags
- [IMPALA-4988](#) - Add query option read_parquet_statistics
- [IMPALA-4996](#) - Single-threaded KuduScanNode
- [IMPALA-5003](#) - Constant propagation in scan conjuncts
- [IMPALA-5008](#) - Fix reading stats for TINYINT and SMALLINT
- [IMPALA-5021](#) - Fix count(*) remaining rows overflow in Parquet.
- [IMPALA-5027](#) - addendum - remove unneeded #define
- [IMPALA-5030](#) - Adds support for NVL2() function
- [IMPALA-5034](#) - Update Breakpad to newer version
- [IMPALA-5038](#) - Fix file size regex to include bytes
- [IMPALA-5041](#) - Allow AuthManager::Init() to be called more than once
- [IMPALA-5042](#) - Use a HashSet instead of ArrayList for O(1) look ups
- [IMPALA-5056](#) - Ensure analysis uses 'fresh' catalog after metadata loading
- [IMPALA-5057](#) - Upgrade glog (0.3.4-p2) and gflags
- [IMPALA-5073](#) - Part 1: add option to use mmap() for buffer pool
- [IMPALA-5077](#) - add NUMA and current cpu to CpuInfo
- [IMPALA-5080](#) - OutOfMemory PermGen space
- [IMPALA-5110](#) - Add deb support to dump_breakpad_symbols.py
- [IMPALA-5111](#) - Fix check when creating NOT NULL PK col in Kudu
- [IMPALA-5113](#) - fix dirty unpinned invariant
- [IMPALA-5120](#) - Default to partitioned join when stats are missing
- [IMPALA-5123](#) - Fix ASAN use after free in timezone_db
- [IMPALA-5125](#) - SimplifyConditionalsRule incorrectly handles aggregates
- [IMPALA-5127](#) - Add history_max option
- [IMPALA-5130](#) - fix race in MemTracker::EnableReservationReporting()
- [IMPALA-5137](#) - Support Kudu UNIXTIME_MICROS as Impala TIMESTAMP.
- [IMPALA-5140](#) - improve docs building guidelines

- [IMPALA-5144](#) - Remove sortBy() hint
- [IMPALA-5147](#) - Add the ability to exclude hosts from query execution
- [IMPALA-5154](#) - Handle 'unpartitioned' Kudu tables
- [IMPALA-5158](#) - Part 1: include untracked memory in MemTracker dumps
- [IMPALA-5166](#) - clean up BufferPool counters
- [IMPALA-5167](#) - Reduce the number of Kudu clients created
- [IMPALA-5169](#) - Add support for async pins in buffer pool
- [IMPALA-5171](#) - update RAT excluded files list
- [IMPALA-5172](#) - crash in tcmmalloc::CentralFreeList::FetchFromOneSpans
- [IMPALA-5173](#) - crash with hash join feeding directly into nlj
- [IMPALA-5174](#) - Bump gflags to 2.2.0-p1
- [IMPALA-5180](#) - Don't use non-deterministic exprs in partition pruning
- [IMPALA-5181](#) - Extract PYPI metadata from a webpage
- [IMPALA-5182](#) - Explicitly close connection to impalad on error from shell
- [IMPALA-5184](#) - build fe against both Hive 1 & 2 APIs
- [IMPALA-5187](#) - , IMPALA-5208: Bump Breakpad Version, undo IMPALA-3794
- [IMPALA-5188](#) - Add slot sorting in TupleDescriptor::LayoutEquals()
- [IMPALA-5189](#) - Pin version of setuptools-scm
- [IMPALA-5192](#) - Don't bake MemPool* into IR
- [IMPALA-5197](#) - Erroneous corrupted Parquet file message
- [IMPALA-5198](#) - Error messages are sometimes dropped before reaching client
- [IMPALA-5207](#) - ,IMPALA-5214: distcc fixes
- [IMPALA-5217](#) - KuduTableSink checks null constraints incorrectly
- [IMPALA-5220](#) - memory maintenance cleanup
- [IMPALA-5221](#) - Avoid re-use of stale SASL contexts.
- [IMPALA-5222](#) - don't call Bits::Log2*() functions
- [IMPALA-5223](#) - Add waiting for HBase Zookeeper nodes to retry loop
- [IMPALA-5224](#) - remove defunct codehaus repository
- [IMPALA-5229](#) - huge page-backed buffers with TCMalloc
- [IMPALA-5230](#) - fix non-functional impalad under ASAN
- [IMPALA-5232](#) - Parquet reader error message prints memory address instead of value
- [IMPALA-5235](#) - Initialize resourceProfile_ with a dummy value
- [IMPALA-5238](#) - transfer reservations between trackers
- [IMPALA-5258](#) - Pass CMAKE_BUILD_TYPE to Impala-izo
- [IMPALA-5259](#) - Add REFRESH FUNCTIONS <db> statement
- [IMPALA-5261](#) - Heap use-after-free in HdfsSequenceTableWriter
- [IMPALA-5273](#) - Replace StringCompare with glibc memcmp
- [IMPALA-5282](#) - Handle overflows in computeResourceProfile().
- [IMPALA-5294](#) - Kudu INSERT partitioning fails with constants
- [IMPALA-5301](#) - Set Kudu minicluster memory limit
- [IMPALA-5304](#) - reduce transfer of Parquet decompression buffers
- [IMPALA-5309](#) - Adds TABLESAMPLE clause for HDFS table refs.
- [IMPALA-5318](#) - Generate access events with fully qualified table names
- [IMPALA-5324](#) - Fix version check in EvalDictionaryFilters
- [IMPALA-5325](#) - Do not update totalHdfsBytes_/numHdfsFiles_ on Catalogd
- [IMPALA-5331](#) - Use new libHDFS API to address "Unknown Error 255"
- [IMPALA-5333](#) - Add support for Impala to work with ADLS
- [IMPALA-5338](#) - Fix Kudu timestamp column default values
- [IMPALA-5339](#) - Fix analysis with sort.columns and expr rewrites
- [IMPALA-5340](#) - Query profile displays stale query state

- [IMPALA-5342](#) - Add comments of loaded tables in the response of GetTables
- [IMPALA-5347](#) - Parquet scanner microoptimizations
- [IMPALA-5355](#) - Fix the order of Sentry roles and privileges
- [IMPALA-5357](#) - Fix unixtime to UTC TimestampValue perf
- [IMPALA-5358](#) - Fix repeatable table sample.
- [IMPALA-5363](#) - Reset probe_batch_ after reaching limit
- [IMPALA-5364](#) - Correct title of query locations table
- [IMPALA-5375](#) - Builds on CentOS 6.4 failing with broken python dependencies
- [IMPALA-5377](#) - Impala may crash if given a fragment instance while restarting
- [IMPALA-5378](#) - Disk IO manager needs to understand ADLS
- [IMPALA-5381](#) - Adds DEFAULT_JOIN_DISTRIBUTION_MODE query option.
- [IMPALA-5383](#) - Fix PARQUET_FILE_SIZE option for ADLS
- [IMPALA-5388](#) - Only retry RPC on lost connection in send call
- [IMPALA-5391](#) - remove C++11 from UDF header
- [IMPALA-5411](#) - Avoid log spew from GetRuntimeProfileStr
- [IMPALA-5419](#) - Check for cancellation when building hash tables
- [IMPALA-5424](#) - Ignore errors when removing minidumps folder
- [IMPALA-5426](#) - Update Hive schema script to 1.1.0
- [IMPALA-5432](#) - Remove invalid DCHECK from SetMemLimitExceeded
- [IMPALA-5438](#) - Always eval union const exprs in subplan.
- [IMPALA-5454](#) - Work around template rendering bug in /memz
- [IMPALA-5469](#) - Fix exception when processing catalog update
- [IMPALA-5487](#) - Fix race in RuntimeProfile::toThrift()
- [IMPALA-5537](#) - Retry RPC on some exceptions with SSL connection
- [IMPALA-5558](#) - /IMPALA-5576: Reopen stale client connection
- [IMPALA-5562](#) - Only recomputeMemLayout() if tuple has a layout.
- [IMPALA-3905](#) - Implements HdfsScanner::GetNext() for LZO text scans.
- [IMPALA-5172](#) - fix incorrect cast in call to LZO decompress

Issues Fixed in Impala for CDH 5.11.2

For the full list of fixed issues for all CDH components in CDH 5.11.2, see [Issues Fixed in CDH 5.11.x](#) on page 2151. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-4276](#) - Profile displays non-default query options set by planner
- [IMPALA-4546](#) - Fix Moscow timezone conversion after 2014
- [IMPALA-4631](#) - don't use floating point operations for time unit conversions
- [IMPALA-4716](#) - Expr rewrite causes IllegalStateException
- [IMPALA-4738](#) - STDDEV_SAMP should return NULL for single record input
- [IMPALA-4962](#) - Fix SHOW COLUMN STATS for HS2
- [IMPALA-5021](#) - Fix count(*) remaining rows overflow in Parquet.
- [IMPALA-5056](#) - Ensure analysis uses 'fresh' catalog after metadata loading
- [IMPALA-5154](#) - Handle 'unpartitioned' Kudu tables
- [IMPALA-5172](#) - Buffer overrun for Snappy decompression
- [IMPALA-5187](#) - Bump breakpad version to include the fix for Breakpad #681, re-enable the strict check that was disabled in IMPALA-3794.
- [IMPALA-5189](#) - Pin version of setuptools-scm
- [IMPALA-5197](#) - Erroneous corrupted Parquet file message
- [IMPALA-5198](#) - Error messages are sometimes dropped before reaching client
- [IMPALA-5217](#) - KuduTableSink checks null constraints incorrectly
- [IMPALA-5223](#) - Add waiting for HBase Zookeeper nodes to retry loop
- [IMPALA-5301](#) - Set Kudu minicluster memory limit

- [IMPALA-5318](#) - Generate access events with fully qualified table names
- [IMPALA-5355](#) - Fix the order of Sentry roles and privileges
- [IMPALA-5363](#) - Reset probe_batch_ after reaching limit
- [IMPALA-5419](#) - Check for cancellation when building hash tables
- [IMPALA-5469](#) - Fix exception when processing catalog update
- [IMPALA-5487](#) - Fix race in RuntimeProfile::toThrift()
- [IMPALA-5524](#) - Fixes NPE during planning with DISABLE_UNSAFE_SPILLS=1
- [IMPALA-5554](#) - sorter DCHECK on null column
- [IMPALA-5580](#) - fix Java UDFs that return NULL strings
- [IMPALA-5615](#) - Fix compute incremental stats for general partition exprs
- [IMPALA-5623](#) - Fix lag() on STRING cols to release UDF mem
- [IMPALA-5638](#) - Fix Kudu table set tblproperties inconsistencies
- [IMPALA-5657](#) - Fix a couple of bugs with FunctionCallExpr and IGNORE NULLS
- [IMPALA-5172](#) - fix incorrect cast in call to LZO decompress

Issues Fixed in Impala for CDH 5.11.1

For the full list of fixed issues for all CDH components in CDH 5.11.1, see [Issues Fixed in CDH 5.11.x](#) on page 2151. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-3641](#) - Fix catalogd RPC responses to DROP IF EXISTS.
- [IMPALA-4088](#) - Assign fix values to the minicluster server ports
- [IMPALA-4293](#) - query profile should include error log
- [IMPALA-4544](#) - ASAN should ignore SEGV and leaks
- [IMPALA-4615](#) - Fix create_table.sql command order
- [IMPALA-4733](#) - Change HBase ports to non-ephemeral
- [IMPALA-4787](#) - Optimize APPX_MEDIAN() memory usage
- [IMPALA-4822](#) - Implement dynamic log level changes
- [IMPALA-4899](#) - Fix parquet table writer dictionary leak
- [IMPALA-4902](#) - Copy parameters map in HdfsPartition.toThrift().
- [IMPALA-4998](#) - Fix missing table lock acquisition.
- [IMPALA-5028](#) - Lock table in /catalog_objects endpoint.
- [IMPALA-5055](#) - Fix DCHECK in parquet-column-readers.cc ReadPageHeader()
- [IMPALA-5088](#) - Fix heap buffer overflow
- [IMPALA-5115](#) - Handle status from HdfsTableSink::WriteClusteredRowBatch
- [IMPALA-5145](#) - Do not constant fold null in CastExprs
- [IMPALA-5156](#) - Drop VLOG level passed into Kudu client
- [IMPALA-5186](#) - Handle failed CreateAndOpenScanner() in MT scan.
- [IMPALA-5193](#) - Initialize decompressor before finding first tuple
- [IMPALA-5251](#) - Fix propagation of input exprs' types in 2-phase agg
- [IMPALA-5252](#) - Fix crash in HiveUdfCall::GetStringVal() when mem_limit exceeded
- [IMPALA-5253](#) - Use appropriate transport for StatestoreSubscriber
- [IMPALA-5322](#) - Fix a potential crash in Frontend & Catalog JNI startup

Issues Fixed in Impala for CDH 5.11.0

For the full list of fixed issues for all CDH components in CDH 5.11.0, see [Issues Fixed in CDH 5.11.x](#) on page 2151. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-1430](#) - IMPALA-4878,IMPALA-4879: codegen native UDAs
- [IMPALA-2020](#) - 4915, 4936: Add rounding for decimal casts
- [IMPALA-2020](#) - IMPALA-4809: Codegen support for DECIMAL_V2
- [IMPALA-2605](#) - Omit the sort and mini stress tests
- [IMPALA-4055](#) - Speed up to_date() with custom implementation.

- [IMPALA-4263](#) - Fix wrong omission of agg/analytic hash exchanges.
- [IMPALA-4282](#) - Remove max length check for type strings.
- [IMPALA-4370](#) - Divide and modulo result types for DECIMAL version V2
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog
- [IMPALA-4675](#) - Case-insensitive matching of Parquet fields.
- [IMPALA-4702](#) - Fix command line help for `webserver_private_key_file`
- [IMPALA-4705](#) - IMPALA-4779, IMPALA-4780: Fix some Expr bugs with codegen
- [IMPALA-4725](#) - Query option to control Parquet array resolution.
- [IMPALA-4729](#) - Implement REPLACE()
- [IMPALA-4742](#) - Change `"{}".format()` to `"{0}".format()` for Py 2.6
- [IMPALA-4749](#) - hit DCHECK in sorter with scratch limit
- [IMPALA-4767](#) - Workaround for HIVE-15653 to preserve table stats.
- [IMPALA-4808](#) - old hash join can reference invalid memory
- [IMPALA-4809](#) - Enable support for DECIMAL_V2 in `decimal_casting.py`
- [IMPALA-4810](#) - fix incorrect expr-test decimal types
- [IMPALA-4810](#) - Make DECIMAL expr-test cases table driven
- [IMPALA-4810](#) - Add DECIMAL_V2 query option
- [IMPALA-4813](#) - Round on divide and multiply
- [IMPALA-4821](#) - Update AVG() for DECIMAL_V2
- [IMPALA-4828](#) - Alter Kudu schema outside Impala may crash on read
- [IMPALA-4854](#) - Fix incremental stats with complex types.
- [IMPALA-4916](#) - Fix maintenance of set of item sets in DisjointSet.
- [IMPALA-4929](#) - Safe concurrent access to IR function call graph
- [IMPALA-4933](#) - IMPALA-4931: Simplify SSL initialization on startup
- [IMPALA-4934](#) - Disable Kudu OpenSSL initialization
- [IMPALA-4981](#) - Re-enable spilling with MT_DOP.
- [IMPALA-4995](#) - Fix integer overflow in `TopNNNode::PrepareForOutput`
- [IMPALA-4997](#) - Fix overflows in `Sorter::TupletIterator`
- [IMPALA-5005](#) - Don't allow server to send SASL COMPLETE msg out of order
- [IMPALA-5025](#) - Update binutils to 2.26.1
- [IMPALA-5027](#) - make udf headers buildable externally

Issues Fixed in Impala for CDH 5.10.2

For the full list of fixed issues for all CDH components in CDH 5.10.2, see [Issues Fixed in CDH 5.10.x](#) on page 2168. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-3641](#) - Fix catalogd RPC responses to DROP IF EXISTS.
- [IMPALA-4088](#) - Assign fix values to the minicluster server ports
- [IMPALA-4293](#) - query profile should include error log
- [IMPALA-4544](#) - ASAN should ignore SEGV and leaks
- [IMPALA-4546](#) - Fix Moscow timezone conversion after 2014
- [IMPALA-4615](#) - Fix `create_table.sql` command order
- [IMPALA-4631](#) - avoid DCHECK in `PlanFragementExecutor::Close()`.
- [IMPALA-4716](#) - Expr rewrite causes `IllegalStateException`
- [IMPALA-4722](#) - Disable log caching in `test_scratch_disk`
- [IMPALA-4725](#) - Query option to control Parquet array resolution.
- [IMPALA-4733](#) - Change HBase ports to non-ephemeral
- [IMPALA-4738](#) - `STDDEV_SAMP` should return NULL for single record input
- [IMPALA-4787](#) - Optimize `APPX_MEDIAN()` memory usage
- [IMPALA-4822](#) - Implement dynamic log level changes
- [IMPALA-4899](#) - Fix parquet table writer dictionary leak

- [IMPALA-4902](#) - Copy parameters map in HdfsPartition.toThrift().
- [IMPALA-4920](#) - custom cluster tests: fix generation of py.test options
- [IMPALA-4998](#) - Fix missing table lock acquisition.
- [IMPALA-5021](#) - Fix count(*) remaining rows overflow in Parquet.
- [IMPALA-5028](#) - Lock table in /catalog_objects endpoint.
- [IMPALA-5055](#) - Fix DCHECK in parquet-column-readers.cc ReadPageHeader()
- [IMPALA-5088](#) - Fix heap buffer overflow
- [IMPALA-5115](#) - Handle status from HdfsTableSink::WriteClusteredRowBatch
- [IMPALA-5145](#) - Do not constant fold null in CastExprs
- [IMPALA-5154](#) - Handle 'unpartitioned' Kudu tables
- [IMPALA-5156](#) - Drop VLOG level passed into Kudu client
- [IMPALA-5172](#) - Buffer overrun for Snappy decompression
- [IMPALA-5183](#) - increase write wait timeout in BufferedBlockMgrTest
- [IMPALA-5186](#) - Handle failed CreateAndOpenScanner() in MT scan.
- [IMPALA-5189](#) - Pin version of setuptools-scm
- [IMPALA-5193](#) - Initialize decompressor before finding first tuple
- [IMPALA-5197](#) - Erroneous corrupted Parquet file message
- [IMPALA-5198](#) - Error messages are sometimes dropped before reaching client
- [IMPALA-5208](#) - Bump toolchain to include fixes for and IMPALA-5187
- [IMPALA-5217](#) - KuduTableSink checks null constraints incorrectly
- [IMPALA-5244](#) - test_hdfs_file_open_fail fails on local filesystem build
- [IMPALA-5252](#) - Fix crash in HiveUdfCall::GetStringVal() when mem_limit exceeded
- [IMPALA-5253](#) - Use appropriate transport for StatestoreSubscriber
- [IMPALA-5287](#) - Test skip.header.line.count on gzip
- [IMPALA-5297](#) - Reduce free-pool-test mem requirement to avoid OOM
- [IMPALA-5301](#) - Set Kudu minicluster memory limit
- [IMPALA-5318](#) - Generate access events with fully qualified table names
- [IMPALA-5322](#) - Fix a potential crash in Frontend & Catalog JNI startup
- [IMPALA-5487](#) - Race in runtime-profile.cc::toThrift() can lead to corrupt profiles being generated while query is running
- [IMPALA-5172](#) - fix incorrect cast in call to LZO decompress

Issues Fixed in Impala for CDH 5.10.1

For the full list of fixed issues for all CDH components in CDH 5.10.1, see [Issues Fixed in CDH 5.10.x](#) on page 2168. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-2605](#) - Omit the sort and mini stress tests.
- [IMPALA-4055](#) - Speed up to_date() with custom implementation.
- [IMPALA-4263](#) - Fix wrong omission of agg/analytic hash exchanges.
- [IMPALA-4282](#) - Remove max length check for type strings.
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog.
- [IMPALA-4675](#) - Case-insensitive matching of Parquet fields.
- [IMPALA-4702](#) - Fix command line help for webserver_private_key_file.
- [IMPALA-4705](#), [IMPALA-4779](#), [IMPALA-4780](#) - Fix some Expr bugs with codegen.
- [IMPALA-4742](#) - Change "{}".format() to "{0}".format() for Py 2.6.
- [IMPALA-4749](#) - hit DCHECK in sorter with scratch limit.
- [IMPALA-4767](#) - Workaround for HIVE-15653 to preserve table stats.
- [IMPALA-4808](#) - old hash join can reference invalid memory.
- [IMPALA-4828](#) - Alter Kudu schema outside Impala may crash on read.
- [IMPALA-4854](#) - Fix incremental stats with complex types.
- [IMPALA-4916](#) - Fix maintenance of set of item sets in DisjointSet.

- [IMPALA-4981](#) - Re-enable spilling with MT_DOP.
- [IMPALA-4995](#) - Fix integer overflow in TopNNNode::PrepareForOutput.
- [IMPALA-4997](#) - Fix overflows in Sorter::TupletIterator.

Issues Fixed in Impala for CDH 5.10.0

For the full list of Impala fixed issues in CDH 5.10 / Impala 2.8, see [this report in the Impala JIRA tracker](#).

For the full list of fixed issues for all CDH components in CDH 5.10.0, see [Issues Fixed in CDH 5.10.x](#) on page 2168. The following list represents the subset of fixed Impala JIRAs from the CDH fixed issues.

- [IMPALA-1169](#) - Admission control info on the queries debug webpage
- [IMPALA-1286](#) - Extract common conjuncts from disjunctions.
- [IMPALA-1430](#) - IMPALA-4108: codegen all builtin aggregate functions
- [IMPALA-1616](#) - Improve the Memory Limit Exceeded error report
- [IMPALA-1654](#) - General partition exprs in DDL operations.
- [IMPALA-1702](#) - Enforce single-table consistency in query analysis
- [IMPALA-1788](#) - Fold constant expressions
- [IMPALA-2013](#) - Reintroduce steps for checking HBase health in run-hbase.sh
- [IMPALA-2057](#) - Better error message for incorrect avro decimal column declaration
- [IMPALA-2521](#) - Add clustered hint to insert statements
- [IMPALA-2523](#) - Make HdfsTableSink aware of clustered input
- [IMPALA-2789](#) - More compact mem layout with null bits at the end.
- [IMPALA-2864](#) - Ensure that client connections are closed after a failed Open()
- [IMPALA-2890](#) - Support ALTER TABLE statements for Kudu tables
- [IMPALA-2905](#) - Move QueryResultSet implementations into separate module
- [IMPALA-2905](#) - Handle coordinator fragment lifecycle like all others
- [IMPALA-2916](#) - Add warning to query profile if debug build
- [IMPALA-2925](#) - Mark test_alloc_update as xfail.
- [IMPALA-3002](#) - IMPALA-1473: Cardinality observability cleanup
- [IMPALA-3125](#) - Fix assignment of equality predicates from an outer-join On-clause
- [IMPALA-3126](#) - Conservative assignment of inner-join On-clause predicates
- [IMPALA-3167](#) - Fix assignment of WHERE conjunct through grouping agg + OJ
- [IMPALA-3200](#) - move bufferpool under runtime
- [IMPALA-3201](#) - in-memory buffer pool implementation
- [IMPALA-3201](#) - reservation implementation for new buffer pool
- [IMPALA-3202](#) - refactor scratch file management into TmpFileMgr
- [IMPALA-3202](#) - DiskIoMgr improvements for new buffer pool
- [IMPALA-3211](#) - provide toolchain build id for bootstrapping
- [IMPALA-3221](#) - Copyright / license audit
- [IMPALA-3229](#) - Don't assume that AUX exists just because of shell env
- [IMPALA-3308](#) - Get expr-test passing on PPC64LE
- [IMPALA-3314](#) - Fix Avro schema loading for partitioned tables
- [IMPALA-3342](#) - Add thread counters to monitor plan fragment execution
- [IMPALA-3346](#) - DeepCopy() Kudu rows into Impala tuples
- [IMPALA-3348](#) - Avoid per-slot check vector size in KuduScanner
- [IMPALA-3398](#) - Add docs to main Impala branch
- [IMPALA-3420](#) - use gold by default
- [IMPALA-3481](#) - Use Kudu ScanToken API for scan ranges
- [IMPALA-3491](#) - Use unique db in test_scanners.py and test_aggregation.py
- [IMPALA-3491](#) - Use unique database fixture in test_partitioning.py
- [IMPALA-3491](#) - Use unique database fixture in test_insert_parquet.py

- [IMPALA-3491](#) - Use unique database fixture in test_nested_types.py
- [IMPALA-3491](#) - Use unique database fixture in test_ddl.py
- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3567](#) - Part 2, IMPALA-3899: factor out PHJ builder
- [IMPALA-3567](#) - move ExecOption profile helpers to RuntimeProfile
- [IMPALA-3586](#) - Clean up union-node.h/cc to enable improvements.
- [IMPALA-3644](#) - Make predicate order deterministic
- [IMPALA-3671](#) - Add query option to limit scratch space usage
- [IMPALA-3676](#) - Use clang as a static analysis tool
- [IMPALA-3710](#) - Kudu DML should ignore conflicts, pt2
- [IMPALA-3710](#) - Kudu DML should ignore conflicts by default
- [IMPALA-3713](#) - ,IMPALA-4439: Fix Kudu DML shell reporting
- [IMPALA-3718](#) - Add test_cancellation tests for Kudu
- [IMPALA-3718](#) - Support subset of functional-query for Kudu
- [IMPALA-3719](#) - Simplify CREATE TABLE statements with Kudu tables
- [IMPALA-3724](#) - Support Kudu non-covering range partitions
- [IMPALA-3725](#) - Support Kudu UPSERT in Impala
- [IMPALA-3726](#) - Add support for Kudu-specific column options
- [IMPALA-3739](#) - Enable stress tests on Kudu
- [IMPALA-3771](#) - Expose kudu client timeout and set default
- [IMPALA-3786](#) - Replace "cloudera" with "apache"
- [IMPALA-3786](#) - Replace "cloudera" with "apache"
- [IMPALA-3788](#) - Fix Kudu ReadMode flag checking
- [IMPALA-3788](#) - Add flag for Kudu read-your-writes
- [IMPALA-3788](#) - Support for Kudu 'read-your-writes' consistency
- [IMPALA-3808](#) - Add incubating DISCLAIMER from the Incubator Branding Guide
- [IMPALA-3809](#) - Show Kudu-specific column metadata in DESCRIBE.
- [IMPALA-3812](#) - Fix error message for unsupported types
- [IMPALA-3815](#) - clean up cross-compiled comparator
- [IMPALA-3823](#) - Add timer to measure Parquet footer reads
- [IMPALA-3838](#) - IMPALA-4495: Codegen EvalRuntimeFilters() and fixes filter stats updates
- [IMPALA-3853](#) - More RAT cleaning.
- [IMPALA-3853](#) - squeasel is MIT (and dual copyright) not Apache
- [IMPALA-3872](#) - allow providing PyPi mirror for python packages
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3884](#) - Support TYPE_TIMESTAMP for HashTableCtx::CodegenAssignNullValue()
- [IMPALA-3902](#) - Scheduler improvements for running multiple fragment instances on a single backend
- [IMPALA-3905](#) - Add single-threaded scan node.
- [IMPALA-3912](#) - test_random_rpc_timeout is flaky.
- [IMPALA-3918](#) - Fix straggler Cloudera -> ASF license headers
- [IMPALA-3920](#) - TotalStorageWaitTime counter not populated for fragments with Kudu scan node
- [IMPALA-3943](#) - Address post-merge comments.
- [IMPALA-3971](#) - IMPALA-3229: Bootstrap an Impala dev environment
- [IMPALA-3973](#) - add position and occurrence to instr()
- [IMPALA-3980](#) - qgen: re-enable Hive as a target database
- [IMPALA-3983](#) - /IMPALA-3974: Delete function jar resources after load
- [IMPALA-4000](#) - Restricted Sentry authorization for Kudu Tables
- [IMPALA-4006](#) - Fix typo in buildall.sh introduced in
- [IMPALA-4006](#) - dangerous rm -rf statements in scripts
- [IMPALA-4008](#) - Don't bake ExprContext pointers into IR code

- [IMPALA-4008](#) - don't bake in hash table and hash join pointers
- [IMPALA-4011](#) - Remove / reword messages when statestore messages are late
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4023](#) - don't attach buffered tuple streams to batches
- [IMPALA-4026](#) - Implement double-buffering for BlockingQueue
- [IMPALA-4028](#) - Improve message for improper Sentry config to make extra spaces visible
- [IMPALA-4037](#) - IMPALA-4038: fix locking during query cancellation
- [IMPALA-4042](#) - Preserve root types when substituting grouping exprs
- [IMPALA-4047](#) - Remove occurrences of 'CDH'/'cdh' from repo
- [IMPALA-4048](#) - Misc. improvements to /sessions
- [IMPALA-4054](#) - Remove serial test workarounds for IMPALA-2479.
- [IMPALA-4056](#) - Fix toSql() of DistributeParam
- [IMPALA-4058](#) - benchmark byteswap on misaligned memory
- [IMPALA-4074](#) - Configuration items duplicate in template of YARN
- [IMPALA-4080](#) - IMPALA-3638: Introduce ExecNode::Codegen()
- [IMPALA-4087](#) - TestFragmentLifecycle.test_failure_in_prepare
- [IMPALA-4091](#) - Fix backend unit to log in logs/be_tests.
- [IMPALA-4096](#) - Allow clean.sh to work from snapshots
- [IMPALA-4097](#) - Crash in kudu-scan-node-test
- [IMPALA-4098](#) - Open()/Close() partition exprs once per fragment instance.
- [IMPALA-4100](#) - 4112: Qgen: Replace EXTRACT UDF + IS [NOT] DISTINCT FROM in HiveSqlWriter
- [IMPALA-4101](#) - qgen: Hive join predicates should only contains equality functions
- [IMPALA-4102](#) - Remote Kudu reads should be reported
- [IMPALA-4104](#) - add DCHECK to ConsumeLocal() and fix tests
- [IMPALA-4110](#) - IMPALA-3853: npm.js uses Artistic License 2
- [IMPALA-4110](#) - Clean up issues found by Apache RAT
- [IMPALA-4110](#) - Apache RAT script on Impala tarballs
- [IMPALA-4111](#) - backend death tests should not produce minidumps
- [IMPALA-4116](#) - Remove 'cdh' from version string again
- [IMPALA-4116](#) - Remove 'cdh' from version string
- [IMPALA-4117](#) - Factor simple scheduler test code into own files
- [IMPALA-4118](#) - extract encryption utils from BufferedBlockMgr
- [IMPALA-4122](#) - qgen: fix bitrotted cluster unit tests
- [IMPALA-4123](#) - Fast bit unpacking
- [IMPALA-4134](#) - IMPALA-3704: Kudu INSERT improvements
- [IMPALA-4136](#) - testKudu planner test hangs if Kudu is not supported
- [IMPALA-4138](#) - Fix AcquireState() for batches that change capacity
- [IMPALA-4142](#) - qgen: Hive does not support CTEs inside sub-query blocks
- [IMPALA-4155](#) - Update default partition when table is altered
- [IMPALA-4160](#) - Remove some leftover Llama references
- [IMPALA-4160](#) - Remove Llama support.
- [IMPALA-4171](#) - Remove JAR from repo.
- [IMPALA-4180](#) - Synchronize accesses to RuntimeState::reader_contexts_
- [IMPALA-4187](#) - Switch RPC latency metrics to histograms
- [IMPALA-4188](#) - Leopard: support external Docker volumes
- [IMPALA-4193](#) - Warn when benchmarks run with sub-optimal CPU settings
- [IMPALA-4194](#) - Bump version to 2.8.0
- [IMPALA-4199](#) - Add 'SNAPSHOT' to Impala version
- [IMPALA-4204](#) - Remove KuduScanNodeTest
- [IMPALA-4205](#) - fix tmp-file-mgr-test under ASAN

- [IMPALA-4206](#) - Add column lineage regression test.
- [IMPALA-4207](#) - test infra: move Hive options from connection to cluster options
- [IMPALA-4213](#) - Fix Kudu predicates that need constant folding
- [IMPALA-4230](#) - ASF policy issues from 2.7.0 rc3.
- [IMPALA-4231](#) - fix codegen time regression
- [IMPALA-4232](#) - qgen: Hive does not support aggregates inside specific analytic clauses
- [IMPALA-4234](#) - Remove astyle config file, looks outdated.
- [IMPALA-4239](#) - fix buffer pool test failures in release build
- [IMPALA-4240](#) - qgen: Add "ParseException line missing)" to Known Errors for Hive
- [IMPALA-4241](#) - remove spurious child queries event
- [IMPALA-4253](#) - impala-server.backends.client-cache.total-clients shows negative value
- [IMPALA-4258](#) - Remove duplicated and unused test macros
- [IMPALA-4259](#) - build Impala without any test cluster setup
- [IMPALA-4260](#) - Alter table add column drops all the column stats
- [IMPALA-4266](#) - Java udf returning string can give incorrect results
- [IMPALA-4269](#) - Codegen merging exchange node
- [IMPALA-4270](#) - Gracefully fail unsupported queries with mt_dop > 0.
- [IMPALA-4274](#) - hang in buffered-block-mgr-test
- [IMPALA-4277](#) - remove unneeded LegacyTCLIService
- [IMPALA-4277](#) - remove references for unsupported s3/s3n connectors
- [IMPALA-4277](#) - allow overriding of Hive/Hadoop versions/locations
- [IMPALA-4278](#) - Don't abort Catalog startup quickly if HMS is not present
- [IMPALA-4283](#) - Ensure Kudu-specific lineage and audit behavior
- [IMPALA-4285](#) - /IMPALA-4286: Fixes for Parquet scanner with MT_DOP > 0.
- [IMPALA-4287](#) - EE tests fail to run when KUDU_IS_SUPPORTED=false
- [IMPALA-4289](#) - Mark agg slots of NDV() functions as non-nullable
- [IMPALA-4291](#) - Reduce LLVM module's preparation time
- [IMPALA-4294](#) - Make check-schema-diff.sh executable from anywhere.
- [IMPALA-4295](#) - XFAIL wildcard SSL test
- [IMPALA-4299](#) - add buildall.sh option to start test cluster
- [IMPALA-4300](#) - Speed up BloomFilter::Or with SIMD
- [IMPALA-4302](#) - ,IMPALA-2379: constant expr arg fixes
- [IMPALA-4303](#) - Do not reset() qualifier of union operands.
- [IMPALA-4309](#) - Introduce Expr rewrite phase and supporting classes
- [IMPALA-4310](#) - Make push_to_asf.py respect --apache_remote
- [IMPALA-4314](#) - Standardize on MT-related data structures
- [IMPALA-4325](#) - StmtRewrite lost parentheses of CompoundPredicate
- [IMPALA-4330](#) - Fix JSON syntax in generate_metrics.py
- [IMPALA-4335](#) - Don't send 0-row batches to clients
- [IMPALA-4338](#) - test infra data migrator: include tables' primary keys in PostgreSQL
- [IMPALA-4339](#) - ensure coredumps end up in IMPALA_HOME
- [IMPALA-4340](#) - explain how to install postgresql-9.5 or higher
- [IMPALA-4343](#) - IMPALA-4354: qgen: model INSERTs; write INSERTs from query model
- [IMPALA-4348](#) - / IMPALA-4333: Improve coordinator fragment cancellation
- [IMPALA-4350](#) - Crash with vlog level 2 in hash join node
- [IMPALA-4352](#) - test infra: store Impala/Kudu primary keys in object model
- [IMPALA-4357](#) - Fix DROP TABLE to pass analysis if the table fails to load
- [IMPALA-4362](#) - Misc. fixes for PFE counters
- [IMPALA-4363](#) - Add Parquet timestamp validation
- [IMPALA-4365](#) - Enabling end-to-end tests on a remote cluster

- [IMPALA-4369](#) - Avoid DCHECK in Parquet scanner with MT_DOP > 0
- [IMPALA-4371](#) - Incorrect DCHECK-s in hdfs-parquet-table-writer
- [IMPALA-4372](#) - 'Describe formatted' returns types in upper case
- [IMPALA-4374](#) - Use new syntax for creating TPC-DS/H tables in Kudu stress test
- [IMPALA-4377](#) - Fix Java UDF-arg buffer use-after-free in UdfExecutorTest.
- [IMPALA-4379](#) - Fix and test Kudu table type checking, follow up
- [IMPALA-4379](#) - Fix and test Kudu table type checking
- [IMPALA-4380](#) - Remove 'cloudera' from hostnames in bin/generate_minidump_collection_testdata.py
- [IMPALA-4381](#) - Incorrect AVX version of BloomFilter::Or
- [IMPALA-4383](#) - Ensure plan fragment report thread is always started
- [IMPALA-4384](#) - NPE when cols list has trailing comma
- [IMPALA-4388](#) - Fix query option reset in tests
- [IMPALA-4391](#) - fix dropped statuses in scanners
- [IMPALA-4392](#) - restore PeakMemoryUsage to DataSink profiles
- [IMPALA-4397](#) - addendum: remove stray semicolon
- [IMPALA-4397](#) - IMPALA-3259: reduce codegen time and memory
- [IMPALA-4403](#) - Implement SHOW RANGE PARTITIONS for Kudu tables
- [IMPALA-4406](#) - Add cryptography export control notice
- [IMPALA-4408](#) - Omit null bytes for Kudu scans with no nullable slots.
- [IMPALA-4409](#) - respect lock order in QueryExecState::CancelInternal()
- [IMPALA-4410](#) - Safer tear-down of RuntimeState
- [IMPALA-4411](#) - Kudu inserts violate lock ordering and could deadlock
- [IMPALA-4412](#) - Per operator timing in profile summary is incorrect when mt_dop > 0
- [IMPALA-4415](#) - Fix unassigned scan range of size 1
- [IMPALA-4421](#) - Send custom cluster & process failure test results to logs/
- [IMPALA-4427](#) - leopard: make DOCKER_IMAGE_NAME required
- [IMPALA-4432](#) - Handle internal codegen disabling properly
- [IMPALA-4433](#) - Always generate testdata using the same time zone setting
- [IMPALA-4434](#) - In Python, ".split('\n') is [''], which has length 1
- [IMPALA-4435](#) - Fix in-predicate-benchmark linking by moving templates
- [IMPALA-4436](#) - StringValue::StringCompare() should match strncmp()
- [IMPALA-4437](#) - fix crash in disk-io-mgr
- [IMPALA-4437](#) - hit DCHECK in buffered-block-mgr-test
- [IMPALA-4438](#) - Serialize test_failpoints.py to reduce memory pressure
- [IMPALA-4440](#) - lineage timestamps can go backwards across daylight savings transitions
- [IMPALA-4441](#) - Divide-by-zero in RuntimeProfile::SummaryStatsCounter::SetStats
- [IMPALA-4442](#) - Fix FE ParserTests UnsatisfiedLinkError
- [IMPALA-4444](#) - Transfer row group resources to row batch on scan failure
- [IMPALA-4446](#) - expr-test fails under ASAN
- [IMPALA-4447](#) - Rein in overly broad sed that dirties the tree
- [IMPALA-4450](#) - qgen: use string concatenation operator for postgres queries
- [IMPALA-4452](#) - Always call AggFnEvaluator::Open() before AggFnEvaluator::Init()
- [IMPALA-4454](#) - test_kudu.TestShowCreateTable flaky
- [IMPALA-4455](#) - MemPoolTest.TryAllocateAligned failure: sizeof v. alignof
- [IMPALA-4458](#) - Fix resource cleanup of cancelled mt scan nodes.
- [IMPALA-4461](#) - Make sure data gets loaded for wide hbase tables.
- [IMPALA-4465](#) - Don't hold process wide lock while serializing Runtime Profile in GetRuntimeProfileStr()
- [IMPALA-4466](#) - Improve Kudu CRUD test coverage
- [IMPALA-4470](#) - Avoid creating a NumericLiteral from NaN/infinity/-0
- [IMPALA-4476](#) - Use unique_database to stop races in test_udfs.py

- [IMPALA-4477](#) - Bump Kudu version to latest master
- [IMPALA-4477](#) - Upgrade Kudu version to latest master
- [IMPALA-4477](#) - Upgrade Kudu version to latest master
- [IMPALA-4477](#) - Upgrade Kudu version to latest master
- [IMPALA-4478](#) - Initial Kudu client mem tracking for sink
- [IMPALA-4479](#) - Use correct isSet() thrift function when evaluating constant bool exprs
- [IMPALA-4480](#) - zero_length_region_ must be as aligned as max_align_t
- [IMPALA-4488](#) - HS2 GetOperationStatus() should keep session alive
- [IMPALA-4490](#) - Only generate runtime filters for hash join nodes.
- [IMPALA-4493](#) - fix string-compare-test when using clang
- [IMPALA-4494](#) - Fix crash in SimpleScheduler
- [IMPALA-4497](#) - Fix Kudu client crash w/ SASL initialization
- [IMPALA-4498](#) - crash in to_utc_timestamp/from_utc_timestamp
- [IMPALA-4502](#) - test_partition_ddl_predicates breaks on non-HDFS filesystems
- [IMPALA-4504](#) - fix races in PlanFragmentExecutor regarding status reporting
- [IMPALA-4509](#) - Initialise Sasl-specific mutex
- [IMPALA-4510](#) - Selectively filter args for metric verification tests
- [IMPALA-4511](#) - Add missing total_time_counter() to PFE::Exec()
- [IMPALA-4512](#) - Add a script that builds Impala on stock Ubuntu 14.04
- [IMPALA-4514](#) - Fix broken exhaustive builds caused by non-nullable columns
- [IMPALA-4516](#) - Don't hold process wide lock connection_to_sessions_map_lock_ while cancelling queries
- [IMPALA-4518](#) - CopyStringVal() doesn't copy null string
- [IMPALA-4519](#) - increase timeout in TestFragmentLifecycle
- [IMPALA-4522](#) - Bound Kudu client threads to avoid stress crash
- [IMPALA-4523](#) - Correct max VARCHAR size to 65535 (2¹⁶ - 1).
- [IMPALA-4525](#) - follow-on: cleanup error handling
- [IMPALA-4525](#) - fix crash when codegen mem limit exceeded
- [IMPALA-4527](#) - Columns in Kudu tables created from Impala default to "NULL"
- [IMPALA-4529](#) - speed up parsing of identifiers
- [IMPALA-4532](#) - Fix use-after-free in ProcessBuildInputAsync()
- [IMPALA-4535](#) - Remove 'auto' from parameter list
- [IMPALA-4539](#) - fix bug when scratch batch references I/O buffers
- [IMPALA-4540](#) - Function call in DCHECK crashes scheduler
- [IMPALA-4541](#) - fix test dimensions for test_codegen_mem_limit
- [IMPALA-4542](#) - Fix use-after-free in some BE tests
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots
- [IMPALA-4553](#) - ntpd must be synchronized for kudu to start.
- [IMPALA-4554](#) - fix projection of nested collections with mt_dop > 0
- [IMPALA-4557](#) - Fix flakiness with FLAGS_stress_free_pool_alloc
- [IMPALA-4561](#) - Replace DISTRIBUTE BY with PARTITION BY in CREATE TABLE
- [IMPALA-4562](#) - Fix for crash on kerberized clusters w/o Kudu support
- [IMPALA-4564](#) - ,IMPALA-4565: mt_dop fixes for old aggs and joins
- [IMPALA-4566](#) - Kudu client glog contention can cause timeouts
- [IMPALA-4567](#) - Fix test_kudu_alter_table exhaustive failures
- [IMPALA-4570](#) - shell tarball breaks with certain setuptools versions
- [IMPALA-4571](#) - Push IN predicates to Kudu
- [IMPALA-4572](#) - Run COMPUTE STATS on Parquet tables with MT_DOP=4
- [IMPALA-4574](#) - Do not treat UUID() like a constant expr
- [IMPALA-4577](#) - Adjust maximum size of row batch queue with MT_DOP
- [IMPALA-4578](#) - Pick up bound predicates for Kudu scan nodes.

- [IMPALA-4579](#) - SHOW CREATE VIEW fails for view containing a subquery
- [IMPALA-4580](#) - Fix crash with FETCH_FIRST when #rows < result cache size
- [IMPALA-4584](#) - Make alter table operations on Kudu tables synchronous
- [IMPALA-4585](#) - Allow the \$DATABASE template in the CATCH section
- [IMPALA-4586](#) - don't constant fold in backend
- [IMPALA-4592](#) - Improve error msg for non-deterministic predicates
- [IMPALA-4594](#) - WriteSlot and CodegenWriteSlot handle escaped NULL slots differently
- [IMPALA-4595](#) - Ignore discarded functions after linking
- [IMPALA-4608](#) - Fix fragment completion times for INSERTs
- [IMPALA-4609](#) - prefix thread counters in fragment profile
- [IMPALA-4613](#) - Make sure timers are finished before sending report profile
- [IMPALA-4614](#) - Set eval cost of timestamp literals.
- [IMPALA-4619](#) - Allow NULL as default value in Kudu tables
- [IMPALA-4628](#) - Disable broken kudu test to unblock GVOs
- [IMPALA-4630](#) - remove debug webpage easter egg
- [IMPALA-4633](#) - Change broken gflag default for Kudu client mem
- [IMPALA-4636](#) - Correct Suse Linux distro string
- [IMPALA-4636](#) - Add support for SLES12 for Kudu integration
- [IMPALA-4638](#) - Run queries with MT_DOP through admission control
- [IMPALA-4642](#) - Fix TestFragmentLifecycle failures; kudu test must wait
- [IMPALA-4654](#) - KuduScanner must return when ReachedLimit()
- [IMPALA-4659](#) - fuzz test fixes
- [IMPALA-4739](#) - ExprRewriter fails on HAVING clauses
- [IMPALA-4765](#) - Avoid using several loading threads on one table
- [IMPALA-4768](#) - Improve logging of table loading
- [IMPALA-3905](#) - Add single-threaded scan node
- [IMPALA-4262](#) - LZO-scanner fails when reading large index files from S3
- [IMPALA-4277](#) - Merge "build against hadoop components in different location" into cdh5-trunk
- [IMPALA-4277](#) - build against hadoop components in different location
- [IMPALA-4322](#) - test_scanners_fuzz.py hits a DCHECK
- [IMPALA-4391](#) - fix dropped status in scanners

Issues Fixed in Impala for CDH 5.9.3

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2194.

- [IMPALA-3794](#) - Workaround for Breakpad ID conflicts
- [IMPALA-4293](#) - query profile should include error log
- [IMPALA-4383](#) - Ensure plan fragment report thread is always started
- [IMPALA-4409](#) - respect lock order in QueryExecState::CancelInternal()
- [IMPALA-4615](#) - Fix create_table.sql command order
- [IMPALA-4787](#) - Optimize APPX_MEDIAN() memory usage
- [IMPALA-5088](#) - Fix heap buffer overflow
- [IMPALA-5193](#) - Initialize decompressor before finding first tuple
- [IMPALA-5197](#) - Erroneous corrupted Parquet file message
- [IMPALA-5252](#) - Fix crash in HiveUdfCall::GetStringVal() when mem_limit exceeded
- [IMPALA-5253](#) - Use appropriate transport for StatestoreSubscriber
- [IMPALA-5355](#) - Fix the order of Sentry roles and privileges
- [IMPALA-5469](#) - Fix exception when processing catalog update

Issues Fixed in Impala for CDH 5.9.2

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2197.

- [IMPALA-1702](#) - Enforce single-table consistency in query analysis.
- [IMPALA-2864](#) - Ensure that client connections are closed after a failed Open()
- [IMPALA-3167](#) - Fix assignment of WHERE conjunct through grouping agg + OJ.
- [IMPALA-3314](#) - Fix Avro schema loading for partitioned tables.
- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3884](#) - Support TYPE_TIMESTAMP for HashTableCtx::CodegenAssignNullValue()
- [IMPALA-3983](#) - [IMPALA-3974](#): Delete function jar resources after load
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4037](#) - [IMPALA-4038](#): fix locking during query cancellation
- [IMPALA-4180](#) - Synchronize accesses to RuntimeState::reader_contexts_
- [IMPALA-4260](#) - Alter table add column drops all the column stats
- [IMPALA-4263](#) - Fix wrong omission of agg/analytic hash exchanges.
- [IMPALA-4266](#) - Java udf returning string can give incorrect results
- [IMPALA-4282](#) - Remove max length check for type strings.
- [IMPALA-4291](#) - Reduce LLVM module's preparation time
- [IMPALA-4363](#) - [IMPALA-4585](#): Add Parquet timestamp validation
- [IMPALA-4391](#) - fix dropped statuses in scanners
- [IMPALA-4433](#) - Always generate testdata using the same time zone setting
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog
- [IMPALA-4488](#) - HS2 GetOperationStatus() should keep session alive
- [IMPALA-4494](#) - [IMPALA-4540](#): Fix crash in SimpleScheduler
- [IMPALA-4516](#) - Don't hold process wide lock connection_to_sessions_map_lock_ while cancelling queries
- [IMPALA-4518](#) - CopyStringVal() doesn't copy null string
- [IMPALA-4539](#) - fix bug when scratch batch references I/O buffers
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots
- [IMPALA-4579](#) - SHOW CREATE VIEW fails for view containing a subquery
- [IMPALA-4705](#) - [IMPALA-4779](#), [IMPALA-4780](#): Fix some Expr bugs with codegen
- [IMPALA-4765](#) - Avoid using several loading threads on one table.
- [IMPALA-4767](#) - Workaround for HIVE-15653 to preserve table stats.
- [IMPALA-4916](#) - Fix maintenance of set of item sets in DisjointSet.
- [IMPALA-4929](#) - Safe concurrent access to IR function call graph
- [IMPALA-4995](#) - Fix integer overflow in TopNNode::PrepareForOutput
- [IMPALA-4997](#) - Fix overflows in Sorter::TupleIterator
- [IMPALA-5005](#) - Don't allow server to send SASL COMPLETE msg out of order
- [IMPALA-4391](#) - fix dropped status in scanners

Issues Fixed in Impala for CDH 5.9.1

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2203.

- [IMPALA-3949](#) - Log error message in FileSystemUtil.copyToLocal()
- [IMPALA-4076](#) - Fix runtime filter sort compare method
- [IMPALA-4099](#) - Fix the error message while loading UDFs with no JARs
- [IMPALA-4120](#) - Incorrect results with LEAD() analytic function
- [IMPALA-4135](#) - Thrift threaded server times out connections during high load
- [IMPALA-4153](#) - Fix count(*) on all blank('') columns - test
- [IMPALA-4170](#) - Fix identifier quoting in COMPUTE INCREMENTAL STATS
- [IMPALA-4196](#) - Cross compile bit-byte-functions
- [IMPALA-4223](#) - Handle truncated file read from HDFS cache
- [IMPALA-4237](#) - Fix materialization of 4-byte decimals in data source scan node
- [IMPALA-4246](#) - SleepForMs() utility function has undefined behavior for > 1s

- [IMPALA-4301](#) - Fix IGNORE NULLS with subquery rewriting
- [IMPALA-4336](#) - Cast exprs after unnesting union operands
- [IMPALA-4387](#) - Validate decimal type in Avro file schema
- [IMPALA-4751](#) - For unknown query IDs, /query_profile_encoded?query_id=123 starts with an empty line.
- [IMPALA-4423](#) - Correct but conservative implementation of Subquery.equals().

Issues Fixed in Impala for CDH 5.9.0

For the full list of Impala fixed issues in CDH 5.9.0 / Impala 2.7.0, see [this report in the Impala JIRA tracker](#).

For the full list of fixed issues for all CDH components in CDH 5.9.0, see [Issues Fixed in CDH 5.9.x](#) on page 2194.

- [IMPALA-1112](#) - Remove some unnecessary code from cross-compilation
- [IMPALA-1240](#) - add back spilling sort now that sorter is not flaky
- [IMPALA-1440](#) - test for insert mem limit
- [IMPALA-3018](#) - Address various small memory allocation related bugs
- [IMPALA-1619](#) - Support 64-bit allocations
- [IMPALA-1633](#) - GetOperationStatus should set errorMessage and sqlState
- [IMPALA-1671](#) - Print time and link to coordinator web UI once query is submitted in shell
- [IMPALA-1683](#) - Allow REFRESH on a single partition
- [IMPALA-2347](#) - Reuse metastore client connections in Catalog
- [IMPALA-2459](#) - Implement next_day date/time UDF
- [IMPALA-2700](#) - ASCII NUL characters are doubled on insert into text tables
- [IMPALA-2767](#) - Web UI call to force expire sessions
- [IMPALA-2878](#) - Fix Base64Decode error and remove duplicate codes
- [IMPALA-2885](#) - ScannerContext::Stream objects should be owned by ScannerContext
- [IMPALA-2979](#) - Fix scheduling on remote hosts
- [IMPALA-3018](#) - Don't return NULL on zero length allocations
- [IMPALA-3063](#) - Separate join inversion from join ordering
- [IMPALA-3084](#) - Cache the sequence of table ref and materialized tuple ids during analysis
- [IMPALA-3181](#) - Add noexcept to some functions
- [IMPALA-3201](#) - buffer pool header only
- [IMPALA-3206](#) - Enable codegen for AVRO_DECIMAL
- [IMPALA-3210](#) - last/first_value() support for IGNORE NULLS
- [IMPALA-3223](#) - Remove boost multiprecision in thirdparty
- [IMPALA-3225](#) - Add script to push from gerrit to ASF
- [IMPALA-3227](#) - generate test TPC data sets during data load
- [IMPALA-3253](#) - Modify gen_build_version.sh to always output the right version
- [IMPALA-3336](#) - qgen: do not randomly generate query options
- [IMPALA-3376](#) - Extra definition level when writing Parquet files
- [IMPALA-3418](#) - The Impala FE project relies on Z-tools snapshot builds
- [IMPALA-3442](#) - Replace '>>' with '>>' in template decls
- [IMPALA-3449](#) - Kudu deploy.py should find clusters by displayName
- [IMPALA-3454](#) - Kudu deletes may fail if subqueries are used
- [IMPALA-3470](#) - DecompressorTest is flaky
- [IMPALA-3491](#) - Merge test_hbase_metadata.py into compute_stats.py. Use unique db fixture
- [IMPALA-3501](#) - ee tests: detect build type and support different timeouts based on the same
- [IMPALA-3507](#) - update binutils version to fix slow linking
- [IMPALA-3521](#) - Impalad should communicate with the statestore after binding to the hs2 and besswax ports
- [IMPALA-3530](#) - Clean up test_ddl.py. Part 1
- [IMPALA-3567](#) - Part 1: groundwork to make Join build sides DataSinks
- [IMPALA-3575](#) - Add retry to backend connection request and rpc timeout

- [IMPALA-3587](#) - Get rid of not_default_fs skip marker
- [IMPALA-3600](#) - Add missing admission control tests
- [IMPALA-3606](#) - Fix Java NPE when trying to add an existing partition
- [IMPALA-3611](#) - track unused Disk IO buffer memory
- [IMPALA-3627](#) - Clean up RPC structures in ImpalaInternalService
- [IMPALA-3632](#) - Add script for running cppclean over the BE code
- [IMPALA-3647](#) - track runtime filter memory in separate tracker
- [IMPALA-3656](#) - Hitting DCHECK/CHECK does not write minidumps
- [IMPALA-3664](#) - S3A test_keys_do_not_work fails
- [IMPALA-3674](#) - Lazy materialization of LLVM module bitcode
- [IMPALA-3677](#) - Write minidump on SIGUSR1
- [IMPALA-3682](#) - Don't retry unrecoverable socket creation errors
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3715](#) - Include total usage of JVM memory
- [IMPALA-3715](#) - Include more info by default in Impala debug memz webpage
- [IMPALA-3716](#) - Add Memory Tab in query's Details page
- [IMPALA-3727](#) - Change microbenchmarks to use percentile-based reporting
- [IMPALA-3729](#) - batch_size=1 coverage for avro scanner
- [IMPALA-3734](#) - C++11 - Replace boost:shared_ptr with std:: equivalent
- [IMPALA-3736](#) - Move Impala HTTP handlers to a separate class
- [IMPALA-3737](#) - Local filesystem build failed loading custom schemas
- [IMPALA-3751](#) - fix clang build errors and warnings
- [IMPALA-3753](#) - Disable create table test for old aggs and joins
- [IMPALA-3756](#) - Fix wrong argument type in HiveStringsTest
- [IMPALA-3757](#) - Add missing lock in RuntimeProfile::ComputeTimeInProfile
- [IMPALA-3762](#) - Download Python requirements before they are needed
- [IMPALA-3763](#) - download_requirements fixes
- [IMPALA-3764](#) - fuzz test HDFS scanners and fix parquet bugs found
- [IMPALA-3767](#) - bootstrap_virtualenv fails to find cython distribution
- [IMPALA-3774](#) - fix download_requirements for older Python versions
- [IMPALA-3778](#) - Fix ASF packaging build
- [IMPALA-3779](#) - Disable cache pool reader thread when HDFS isn't running
- [IMPALA-3780](#) - avoid many small reads past end of block
- [IMPALA-3786](#) - Remove "Cloudera" from impalad webpage title
- [IMPALA-3790](#) - AC tests timeout in codecoverage builds
- [IMPALA-3799](#) - Make MAX_SCAN_RANGE_LENGTH accept formatted quantities
- [IMPALA-3806](#) - remove a few modern shell idioms to improve RHEL5 support
- [IMPALA-3817](#) - Ensure filter hash function is the same on all hardware
- [IMPALA-3839](#) - Fix race condition in impala_cluster.py
- [IMPALA-3843](#) - Update warning for non-SSSE3 CPUs
- [IMPALA-3845](#) - Split up hdfs-parquet-scanner.cc into more files/components
- [IMPALA-3852](#) - Remove Derby and Shiro FE dependencies
- [IMPALA-3854](#) - Fix use-after-free in HdfsTextScanner::Close()
- [IMPALA-3856](#) - Fix BinaryPredicate normalization for Kudu
- [IMPALA-3857](#) - KuduScanNode race on returning "optional" threads
- [IMPALA-3864](#) - qgen: reduce likelihood of create_query() exceptions
- [IMPALA-3866](#) - consistent user-facing terminology for scratch dirs
- [IMPALA-3881](#) - Add DataTables 1.10.12 to www/
- [IMPALA-3886](#) - Improve log of pip_download.py
- [IMPALA-3892](#) - qgen: always run Impala with -convert_legacy_hive_parquet_utc_timestamps=true

- [IMPALA-3905](#) - Add HdfsScanner::GetNext() interface and implementation for Parquet
- [IMPALA-3906](#) - Materialize implicitly referenced IR functions
- [IMPALA-3914](#) - SKIP_TOOLCHAIN_BOOTSTRAP skips Python package downloads
- [IMPALA-3918](#) - Remove Cloudera copyrights and add ASF license header
- [IMPALA-3923](#) - fix overflow in BufferedTupleStream::GetRows()
- [IMPALA-3924](#) - Ubuntu16 support
- [IMPALA-3936](#) - BufferedBlockMgr fixes for Pin() while write in flight
- [IMPALA-3939](#) - Data loading may fail on tpch kudu
- [IMPALA-3943](#) - Do not throw scan errors for empty Parquet files
- [IMPALA-3946](#) - fix MemPool integrity issues with empty chunks
- [IMPALA-3952](#) - Clear scratch batch mem pool if Open() failed
- [IMPALA-3953](#) - Fixes for KuduScanNode BE test failure
- [IMPALA-3954](#) - Add unique_database to scanner test
- [IMPALA-3957](#) - Test failure in S3 build: TestLoadData.test_load
- [IMPALA-3964](#) - Fix crash when a count(*) is performed on a nested collection
- [IMPALA-3969](#) - stress test: add option to set common query options
- [IMPALA-3972](#) - Improve display of /varz page
- [IMPALA-3992](#) - bad shell error message when running nonexistent file

Issues Fixed in Impala for CDH 5.8.5

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2232.

- [IMPALA-1346](#) - /1590/2344: fix sorter buffer mgmt when spilling
- [IMPALA-1619](#), [IMPALA-3018](#): Address various small memory allocation related bugs
- [IMPALA-1619](#) - Support 64-bit allocations.
- [IMPALA-1657](#) - Rework detection and reporting of corrupt table stats.
- [IMPALA-2864](#) - Ensure that client connections are closed after a failed Open()
- [IMPALA-3018](#) - Don't return NULL on zero length allocations.
- [IMPALA-3159](#) - impala-shell does not accept wildcard or SAN certificates
- [IMPALA-3167](#) - Fix assignment of WHERE conjunct through grouping agg + OJ.
- [IMPALA-3314](#) - Fix Avro schema loading for partitioned tables.
- [IMPALA-3344](#) - Simplify sorter and document/enforce invariants.
- [IMPALA-3441](#), [IMPALA-3659](#): check for malformed Avro data
- [IMPALA-3499](#) - Split catalog update
- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3575](#) - Add retry to backend connection request and rpc timeout
- [IMPALA-3628](#) - Fix cancellation from shell when security is enabled
- [IMPALA-3633](#) - cancel fragment if coordinator is gone
- [IMPALA-3646](#) - Handle corrupt RLE literal or repeat counts of 0.
- [IMPALA-3670](#) - fix sorter buffer mgmt bugs
- [IMPALA-3678](#) - Fix migration of predicates into union operands with an order by + limit.
- [IMPALA-3680](#) - Cleanup the scan range state after failed hdfs cache reads
- [IMPALA-3682](#) - Don't retry unrecoverable socket creation errors
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3711](#) - Remove unnecessary privilege checks in getDbsMetadata()
- [IMPALA-3732](#) - handle string length overflow in avro files
- [IMPALA-3745](#) - parquet invalid data handling
- [IMPALA-3751](#) - fix clang build errors and warnings
- [IMPALA-3754](#) - fix TestParquet.test_corrupt_rle_counts flakiness
- [IMPALA-3776](#) - fix 'describe formatted' for Avro tables
- [IMPALA-3820](#) - Handle linkage errors while loading Java UDFs in Catalog

- [IMPALA-3861](#) - Replace BetweenPredicates with their equivalent CompoundPredicate.
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3884](#) - Support TYPE_TIMESTAMP for HashTableCtx::CodegenAssignNullValue()
- [IMPALA-3915](#) - Register privilege and audit requests when analyzing resolved table refs.
- [IMPALA-3930](#), [IMPALA-2570](#): Fix shuffle insert hint with constant partition exprs.
- [IMPALA-3940](#) - Fix getting column stats through views.
- [IMPALA-3949](#) - Log the error message in FileSystemUtil.copyToLocal()
- [IMPALA-3964](#) - Fix crash when a count(*) is performed on a nested collection.
- [IMPALA-3965](#) - TSSLSocketWithWildcardSAN.py not exported as part of impala-shell build lib
- [IMPALA-3983](#), [IMPALA-3974](#): Delete function jar resources after load
- [IMPALA-4019](#) - initialize member variables in HdfsTableSink
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4037](#), [IMPALA-4038](#): fix locking during query cancellation
- [IMPALA-4049](#) - fix empty batch handling NLJ build side
- [IMPALA-4076](#) - Fix runtime filter sort compare method
- [IMPALA-4099](#) - Fix the error message while loading UDFs with no JARs
- [IMPALA-4120](#) - Incorrect results with LEAD() analytic function
- [IMPALA-4135](#) - Thrift threaded server times-out connections during high load
- [IMPALA-4153](#) - Fix count(*) on all blank("") columns - test
- [IMPALA-4170](#) - Fix identifier quoting in COMPUTE INCREMENTAL STATS.
- [IMPALA-4180](#) - Synchronize accesses to RuntimeState::reader_contexts_
- [IMPALA-4196](#) - Cross compile bit-byte-functions
- [IMPALA-4223](#) - Handle truncated file read from HDFS cache
- [IMPALA-4237](#) - Fix materialization of 4 byte decimals in data source scan node.
- [IMPALA-4246](#) - SleepForMs() utility function has undefined behavior for > 1s
- [IMPALA-4260](#) - Alter table add column drops all the column stats
- [IMPALA-4263](#) - Fix wrong omission of agg/analytic hash exchanges.
- [IMPALA-4266](#) - Java udf returning string can give incorrect results
- [IMPALA-4282](#) - Remove max length check for type strings.
- [IMPALA-4293](#) - query profile should include error log
- [IMPALA-4295](#) - XFAIL wildcard SSL test
- [IMPALA-4336](#) - Cast exprs after unnesting union operands.
- [IMPALA-4363](#) - Add Parquet timestamp validation
- [IMPALA-4383](#) - Ensure plan fragment report thread is always started
- [IMPALA-4391](#) - fix dropped statuses in scanners
- [IMPALA-4423](#) - Correct but conservative implementation of Subquery.equals().
- [IMPALA-4433](#) - Always generate testdata using the same time zone setting
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog This commit fixes an issue where multiple long-running operations on the same catalog object (e.g. table) can block other catalog operations from making progress.
- [IMPALA-4488](#) - HS2 GetOperationStatus() should keep session alive
- [IMPALA-4518](#) - CopyStringVal() doesn't copy null string
- [IMPALA-4539](#) - fix bug when scratch batch references I/O buffers
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots
- [IMPALA-4579](#) - SHOW CREATE VIEW fails for view containing a subquery
- [IMPALA-4765](#) - Avoid using several loading threads on one table.
- [IMPALA-4767](#) - Workaround for HIVE-15653 to preserve table stats.
- [IMPALA-4779](#), [IMPALA-4780](#): Fix conditional functions built-in and Timestamp bounds
- [IMPALA-4787](#) - Optimize APPX_MEDIAN() memory usage
- [IMPALA-4916](#) - Fix maintenance of set of item sets in DisjointSet.
- [IMPALA-4995](#) - Fix integer overflow in TopNNNode::PrepareForOutput

- [IMPALA-4997](#) - Fix overflows in Sorter::TupleIterator
- [IMPALA-5005](#) - Don't allow server to send SASL COMPLETE msg out of order
- [IMPALA-5088](#) - Fix heap buffer overflow
- [IMPALA-5253](#) - Use appropriate transport for StatestoreSubscriber

Issues Fixed in Impala for CDH 5.8.4

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2247.

- [IMPALA-1702](#) - "invalidate metadata" can cause duplicate TableIds
- [IMPALA-3167](#) - Fix assignment of WHERE clause predicate through grouping aggregate and outer join
- [IMPALA-3314](#) - Fix Avro schema loading for partitioned tables
- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3575](#) - Add retry to backend connection request and rpc timeout
- [IMPALA-3682](#) - Don't retry unrecoverable socket creation errors
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3884](#) - Support TYPE_TIMESTAMP for HashTableCtx::CodegenAssignNullValue()
- [IMPALA-3949](#) - Log the error message in FileSystemUtil.copyToLocal()
- [IMPALA-3964](#) - Fix crash when a count(*) is performed on a nested collection.
- [IMPALA-3983](#) - Delete function jar resources after load
- [IMPALA-4037](#) - ChildQuery::Cancel() appears to violate lock ordering
- [IMPALA-4038](#) - Fix locking during query cancellation
- [IMPALA-4076](#) - Fix runtime filter sort compare method
- [IMPALA-4099](#) - Fix the error message while loading UDFs with no JARs
- [IMPALA-4120](#) - Incorrect results with LEAD() analytic function
- [IMPALA-4153](#) - Fix count(*) on all blank("") columns - test
- [IMPALA-4223](#) - Handle truncated file read from HDFS cache
- [IMPALA-4246](#) - SleepForMs() utility function has undefined behavior for > 1s
- [IMPALA-4336](#) - Cast expressions after unnesting union operands
- [IMPALA-4363](#) - Add Parquet timestamp validation
- [IMPALA-4391](#) - Fix dropped statuses in scanners
- [IMPALA-4423](#) - Correct but conservative implementation of Subquery.equals()
- [IMPALA-4433](#) - Always generate test data using the same time zone setting
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog. Fixes an issue where multiple long-running operations on the same catalog object (for example, a table) can block other catalog operations from making progress
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots
- [IMPALA-4391](#) - Fix dropped status in scanners

Issues Fixed in Impala for CDH 5.8.3

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2251.

- [IMPALA-1619](#) - Support 64-bit allocations
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3751](#) - Fix clang build errors and warnings
- [IMPALA-4135](#) - Thrift threaded server times-out connections during high load
- [IMPALA-4170](#) - Fix identifier quoting in COMPUTE INCREMENTAL STATS
- [IMPALA-4180](#) - Synchronize accesses to RuntimeState::reader_contexts_
- [IMPALA-4196](#) - Cross compile bit-byte-functions
- [IMPALA-4237](#) - Fix materialization of 4-byte decimals in data source scan node

Issues Fixed in Impala for CDH 5.8.2

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2253.

- [IMPALA-1346](#) - /1590/2344: fix sorter buffer mgmt when spilling
- [IMPALA-3159](#) - impala-shell does not accept wildcard or SAN certificates
- [IMPALA-3344](#) - Simplify sorter and document/enforce invariants.
- [IMPALA-3441](#) - , IMPALA-3659: check for malformed Avro data
- [IMPALA-3499](#) - Split catalog update.
- [IMPALA-3628](#) - Fix cancellation from shell when security is enabled
- [IMPALA-3633](#) - cancel fragment if coordinator is gone
- [IMPALA-3646](#) - Handle corrupt RLE literal or repeat counts of 0.
- [IMPALA-3670](#) - fix sorter buffer mgmt bugs
- [IMPALA-3678](#) - Fix migration of predicates into union operands with an order by + limit.
- [IMPALA-3680](#) - Cleanup the scan range state after failed hdfs cache reads
- [IMPALA-3711](#) - Remove unnecessary privilege checks in getDbsMetadata().
- [IMPALA-3732](#) - handle string length overflow in avro files
- [IMPALA-3745](#) - parquet invalid data handling
- [IMPALA-3754](#) - fix TestParquet.test_corrupt_rle_counts flakiness
- [IMPALA-3772](#) - Fix admission control stress test.
- [IMPALA-3776](#) - fix 'describe formatted' for Avro tables
- [IMPALA-3820](#) - Handle linkage errors while loading Java UDFs in Catalog
- [IMPALA-3861](#) - Replace BetweenPredicates with their equivalent CompoundPredicate.
- [IMPALA-3915](#) - Register privilege and audit requests when analyzing resolved table refs.
- [IMPALA-3930](#) - Fix shuffle insert hint with constant partition exprs.
- [IMPALA-3940](#) - Fix getting column stats through views.
- [IMPALA-3965](#) - TSSLSocketWithWildcardSAN.py not exported as part of impala-shell build lib
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4049](#) - fix empty batch handling NLJ build side

Issues Fixed in Impala for CDH 5.8.0

The following list contains the most critical fixed issues (`priority='Blocker'`) from the JIRA system. For the full list of fixed issues in CDH 5.8.0 / Impala 2.6.0, see [this report in the Impala JIRA tracker](#).

RuntimeState::error_log_ crashes

A crash could occur, with stack trace pointing to `impala::RuntimeState::ErrorLog`.

Bug: [IMPALA-3385](#)

Severity: High

`HiveUdfCall::Open()` produces unsynchronized access to `JniUtil::global_refs_vector`

A crash could occur because of contention between multiple calls to Java UDFs.

Bug: [IMPALA-3378](#)

Severity: High

`HBaseTableWriter::CreatePutList()` produces unsynchronized access to `JniUtil::global_refs_vector`

A crash could occur because of contention between multiple concurrent statements writing to HBase.

Bug: [IMPALA-3379](#)

Severity: High

Stress test failure: `sorter.cc:745] Check failed: i == 0 (1 vs. 0)`

A crash or wrong results could occur if the spill-to-disk mechanism encountered a zero-length string at the very end of a data block.

Bug: [IMPALA-3317](#)

Severity: High

String data coming out of agg can be corrupted by blocking operators

If a query plan contains an aggregation node producing string values anywhere within a subplan (that is, if in the SQL statement, the aggregate function appears within an inline view over a collection column), the results of the aggregation may be incorrect.

Bug: [IMPALA-3311](#)

Severity: High

CTAS with subquery throws AuthzException

A `CREATE TABLE AS SELECT` operation could fail with an authorization error, due to a slight difference in the privilege checking for the CTAS operation.

Bug: [IMPALA-3269](#)

Severity: High

Crash on inserting into table with binary and parquet

Impala incorrectly allowed `BINARY` to be specified as a column type, resulting in a crash during a write to a Parquet table with a column of that type.

Bug: [IMPALA-3237](#)

Severity: High

`RowBatch::MaxTupleBufferSize()` calculation incorrect, may lead to memory corruption

A crash could occur while querying tables with very large rows, for example wide tables with many columns or very large string values. This problem was identified in Impala 2.3, but had low reproducibility in subsequent releases. The fix ensures the memory allocation size is correct.

Bug: [IMPALA-3105](#)

Severity: High

Thrift buffer overflows when serialize more than 3355443200 bytes in impala

A very large memory allocation within the `catalogd` daemon could exceed an internal Thrift limit, causing a crash.

Bug: [IMPALA-3494](#)

Severity: High

Altering table partition's storage format is not working and crashing the daemon

If a partitioned table used a file format other than Avro, and the file format of an individual partition was changed to Avro, subsequent queries could encounter a crash.

Bug: [IMPALA-3314](#)

Severity: High

Race condition may cause scanners to spin with runtime filters on Avro or Sequence files

A timing problem during runtime filter processing could cause queries against Avro or SequenceFile tables to hang.

Bug: [IMPALA-3798](#)

Severity: High

Issues Fixed in Impala for CDH 5.7.6

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2264.

- [IMPALA-2864](#) - Ensure that client connections are closed after a failed `Open()`
- [IMPALA-3167](#) - Fix assignment of WHERE conjunct through grouping agg + OJ.

- [IMPALA-3552](#) - Make incremental stats max serialized size configurable
- [IMPALA-3698](#) - Fix Isilon permissions test
- [IMPALA-3861](#) - Replace BetweenPredicates with their equivalent CompoundPredicate.
- [IMPALA-3875](#) - Thrift threaded server hang in some cases
- [IMPALA-3983](#) - /IMPALA-3974: Delete function jar resources after load
- [IMPALA-4153](#) - Return valid non-NULL pointer for 0-byte allocations
- [IMPALA-4223](#) - Handle truncated file read from HDFS cache
- [IMPALA-4336](#) - Cast exprs after unnesting union operands.
- [IMPALA-4363](#) - Add Parquet timestamp validation
- [IMPALA-4423](#) - Correct but conservative implementation of Subquery.equals().
- [IMPALA-4433](#) - Always generate testdata using the same time zone setting
- [IMPALA-4449](#) - Revisit table locking pattern in the catalog
- [IMPALA-4488](#) - HS2 GetOperationStatus() should keep session alive
- [IMPALA-4550](#) - Fix CastExpr analysis for substituted slots
- [IMPALA-4579](#) - SHOW CREATE VIEW fails for view containing a subquery
- [IMPALA-4765](#) - Avoid using several loading threads on one table.

Issues Fixed in Impala for CDH 5.7.5

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2267.

- [IMPALA-1619](#) - Support 64-bit allocations
- [IMPALA-1740](#) - Add support for skip.header.line.count
- [IMPALA-3458](#) - Fix table creation to test insert with header lines
- [IMPALA-3949](#) - Log the error message in FileSystemUtil.copyToLocal()
- [IMPALA-4037](#) - Fix locking during query cancellation
- [IMPALA-4076](#) - Fix runtime filter sort compare method
- [IMPALA-4099](#) - Fix the error message while loading UDFs with no JARs
- [IMPALA-4120](#) - Incorrect results with LEAD() analytic function
- [IMPALA-4135](#) - Thrift threaded server times-out connections during high load
- [IMPALA-4170](#) - Fix identifier quoting in COMPUTE INCREMENTAL STATS
- [IMPALA-4196](#) - Cross compile bit-byte functions
- [IMPALA-4237](#) - Fix materialization of 4 byte decimals in data source scan node
- [IMPALA-4246](#) - SleepForMs() utility function has undefined behavior for > 1s

Issues Fixed in Impala for CDH 5.7.4

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2269.

- [IMPALA-3081](#) - Increase memory limit for TestWideRow
- [IMPALA-3311](#) - Fix string data coming out of aggs in subplans
- [IMPALA-3575](#) - Add retry to back end connection request and rpc timeout
- [IMPALA-3678](#) - Fix migration of predicates into union operands with an order by + limit.
- [IMPALA-3682](#) - Do not retry unrecoverable socket creation errors
- [IMPALA-3687](#) - Fix test failure introduced by backporting
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3820](#) - Handle linkage errors while loading Java UDFs in Catalog
- [IMPALA-3930](#) - Fix shuffle insert hint with constant partition exprs
- [IMPALA-3940](#) - Fix getting column stats through views
- [IMPALA-4020](#) - Handle external conflicting changes to HMS gracefully
- [IMPALA-4049](#) - Fix empty batch handling NLJ build side

Issues Fixed in Impala for CDH 5.7.2

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2274.

- [IMPALA-1928](#) - Fix Thrift client transport wrapping order
- [IMPALA-2660](#) - Respect auth_to_local configs from hdfs configs
- [IMPALA-3276](#) - Consistently handle pin failure in BTS::PrepareForRead()
- [IMPALA-3369](#) - Add ALTER TABLE SET COLUMN STATS statement.
- [IMPALA-3441](#) - Impala should not crash for invalid avro serialized data
- [IMPALA-3499](#) - Split catalog update
- [IMPALA-3502](#) - Fix race in the coordinator while updating filter routing table
- [IMPALA-3633](#) - Cancel fragment if coordinator is gone
- [IMPALA-3732](#) - Handle string length overflow in Avro files
- [IMPALA-3745](#) - Corrupt encoded values in parquet files can cause crashes
- [IMPALA-3751](#) - Fix clang build errors and warnings
- [IMPALA-3754](#) - Fix TestParquet.test_corrupt_rle_counts flakiness

Issues Fixed in Impala for CDH 5.7.1

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2277.

- [IMPALA-2076](#) - Correct execution time tracking for DataStreamSender
- [IMPALA-2502](#) - Don't redundantly repartition grouping aggregations
- [IMPALA-2892](#) - Buffered-tuple-stream-ir.cc is not cross-compiled
- [IMPALA-3133](#) - Wrong privileges after a REVOKE ALL ON SERVER statement
- [IMPALA-3139](#) - Fix drop table statement to not drop views and vice versa
- [IMPALA-3141](#) - Send dummy filters when filter production is disabled
- [IMPALA-3194](#) - Allow queries materializing scalar type columns in RC/sequence files
- [IMPALA-3220](#) - Skip logging empty ScannerContext's stream in parse error
- [IMPALA-3236](#) - Increase timeout for runtime filter tests
- [IMPALA-3238](#) - Avoid log spam for very large hash tables
- [IMPALA-3245](#), [IMPALA-3305](#): Fix crash with global filters when NUM_NODES=1
- [IMPALA-3269](#) - Remove authz checks on default table location in CTAS queries
- [IMPALA-3285](#) - Fix ASAN failure in webserver-test
- [IMPALA-3317](#) - Fix crash in sorter when spilling zero-length strings
- [IMPALA-3334](#) - Fix some bugs in query options' parsing.
- [IMPALA-3367](#) - Ensure runtime filters tests run on 3 nodes
- [IMPALA-3378](#), [IMPALA-3379](#): fix various JNI issues
- [IMPALA-3385](#) - Fix crashes on accessing error_log
- [IMPALA-3395](#) - Old HT filter code uses wrong expr type
- [IMPALA-3396](#) - Fix ConcurrentTimerCounter unit test "TimerCounterTest" failure.
- [IMPALA-3412](#) - Fix CHAR codegen crash in tuple comparator
- [IMPALA-3420](#) - Set IMPALA_THRIFT_VERSION patch level to +4

Issues Fixed in Impala for CDH 5.7.0

The following list contains the most critical issues (`priority='Blocker'`) from the JIRA system. For the full list of fixed issues in CDH 5.7.0 / Impala 2.5.0, see [this report in the Impala JIRA tracker](#).

Stress test hit assert in LLVM: external function could not be resolved

Bug: [IMPALA-2683](#)

The stress test was running a build with the TPC-H, TPC-DS, and TPC-H nested queries with scale factor 3.

Impalad is crashing if udf jar is not available in hdfs location for first time

Bug: [IMPALA-2365](#)

If a UDF JAR was not available in the HDFS location specified in the `CREATE FUNCTION` statement, the `impalad` daemon could crash.

PAGG hits `mem_limit` when switching to I/O buffers

Bug: [IMPALA-2535](#)

A join query could fail with an out-of-memory error despite the apparent presence of sufficient memory. The cause was the internal ordering of operations that could cause a later phase of the query to allocate memory required by an earlier phase of the query. The workaround was to either increase or decrease the `MEM_LIMIT` query option, because the issue would only occur for a specific combination of memory limit and data volume.

Prevent migrating incorrectly inferred identity predicates into inline views

Bug: [IMPALA-2643](#)

Referring to the same column twice in a view definition could cause the view to omit rows where that column contained a `NULL` value. This could cause incorrect results due to an inaccurate `COUNT(*)` value or rows missing from the result set.

Fix migration/assignment of `ON`-clause predicates inside inline views

Bug: [IMPALA-1459](#)

Some combinations of `ON` clauses in join queries could result in comparisons being applied at the wrong stage of query processing, leading to incorrect results. Wrong predicate assignment could happen under the following conditions:

- The query includes an inline view that contains an outer join.
- That inline view is joined with another table in the enclosing query block.
- That join has an `ON` clause containing a predicate that only references columns originating from the outer-joined tables inside the inline view.

Wrong plan of `NOT IN` aggregate subquery when a constant is used in subquery predicate

Bug: [IMPALA-2093](#)

`IN` subqueries might return wrong results if the left-hand side of the `IN` is a constant. For example:

```
select * from alltypestiny t1
  where 10 not in (select sum(int_col) from alltypestiny);
```

Parquet DictDecoders accumulate throughout query

Bug: [IMPALA-2940](#)

Parquet dictionary decoders can accumulate throughout query execution, leading to excessive memory usage. One decoder is created per-column per-split.

Planner doesn't set the `has_local_target` field correctly

Bug: [IMPALA-3056](#)

MemPool allocation growth behavior

Bug: [IMPALA-2742](#)

Currently, the MemPool would always double the size of the last allocation. This can lead to bad behavior if the MemPool transferred the ownership of all its data except the last chunk. In the next allocation, the next allocated chunk would double the size of this large chunk, which can be undesirable.

Drop partition operations don't follow the catalog's locking protocol

Bug: [IMPALA-3035](#)

The `CatalogOpExecutor.alterTableDropPartition()` function violates the locking protocol used in the catalog that requires `catalogLock_` to be acquired before any table-level lock. That may cause deadlocks when `ALTER TABLE DROP PARTITION` is executed concurrently with other DDL operations.

HAVING clause without aggregation not applied properly

Bug: [IMPALA-2215](#)

A query with a `HAVING` clause but no `GROUP BY` clause was not being rejected, despite being invalid syntax. For example:

```
select case when 1=1 then 'didit' end as c1 from (select 1 as one) a having 1!=1;
```

Hit DCHECK Check failed: HasDateOrTime()

Bug: [IMPALA-2914](#)

`TimestampValue::ToTimestampVal()` requires a valid `TimestampValue` as input. This requirement was not enforced in some places, leading to serious errors.

Aggregation spill loop gives up too early leading to mem limit exceeded errors

Bug: [IMPALA-2986](#)

An aggregation query could fail with an out-of-memory error, despite sufficient memory being reported as available.

`DataStreamSender::Channel::CloseInternal()` does not close the channel on an error.

Bug: [IMPALA-2592](#)

Some queries do not close an internal communication channel on an error. This will cause the node on the other side of the channel to wait indefinitely, causing the query to hang. For example, this issue could happen on a Kerberos-enabled system if the credential cache was outdated. Although the affected query hangs, the `impalad` daemons continue processing other queries.

Codegen does not catch exceptions in FROM_UNIXTIME()

Bug: [IMPALA-2184](#)

Querying for the min or max value of a timestamp cast from a bigint via `from_unixtime()` fails silently and crashes instances of `impalad` when the input includes a value outside of the valid range.

Workaround: Disable native code generation with:

```
SET disable_codegen=true;
```

Impala returns wrong result for function 'conv(bigint, from_base, to_base)'

Bug: [IMPALA-2788](#)

Impala returns wrong result for function `conv()`. Function `conv(bigint, from_base, to_base)` returns an correct result, while `conv(string, from_base, to_base)` returns the correct value. For example:

```
select 2061013007, conv(2061013007, 16, 10), conv('2061013007', 16, 10);
+-----+-----+-----+
| 2061013007 | conv(2061013007, 16, 10) | conv('2061013007', 16, 10) |
+-----+-----+-----+
| 2061013007 | 1627467783                | 139066421255                |
+-----+-----+-----+
Fetched 1 row(s) in 0.65s
```

```
select 2061013007, conv(cast(2061013007 as bigint), 16, 10), conv('2061013007', 16, 10);
+-----+-----+-----+
| 2061013007 | conv(cast(2061013007 as bigint), 16, 10) | conv('2061013007', 16, 10) |
+-----+-----+-----+
```

```

+-----+-----+-----+
| 2061013007 | 1627467783 | 139066421255 |
+-----+-----+-----+

select 2061013007, conv(cast(2061013007 as string), 16, 10), conv('2061013007', 16, 10);
+-----+-----+-----+
| 2061013007 | conv(cast(2061013007 as string), 16, 10) | conv('2061013007', 16, 10) |
+-----+-----+-----+
| 2061013007 | 139066421255 | 139066421255 |
+-----+-----+-----+

select 2061013007, conv(cast(cast(2061013007 as decimal(20,0)) as bigint), 16, 10),
conv('2061013007', 16, 10);
+-----+-----+-----+
| 2061013007 | conv(cast(cast(2061013007 as decimal(20,0)) as bigint), 16, 10) |
conv('2061013007', 16, 10) |
+-----+-----+-----+
| 2061013007 | 1627467783 | 139066421255 |
+-----+-----+-----+

```

Workaround: Cast the value to string and use `conv(string, from_base, to_base)` for conversion.

Issues Fixed in Impala for CDH 5.6.1

For the full list of fixed issues for all CDH components, see [Upstream Issues Fixed](#) on page 2291.

- [IMPALA-852](#), [IMPALA-2215](#) - Analyze HAVING clause before aggregation
- [IMPALA-1092](#) - Fix estimates for trivial coord-only queries
- [IMPALA-1170](#) - Fix URL parsing when path contains '@'
- [IMPALA-1934](#) - Allow shell to retrieve LDAP password from shell cmd
- [IMPALA-2093](#) - Disallow NOT IN aggregate subqueries with a constant lhs expr
- [IMPALA-2184](#) - don't inline timestamp methods with try/catch blocks in IR
- [IMPALA-2425](#) - Broadcast join hint not enforced when low memory limit is set
- [IMPALA-2503](#) - Add missing String.format() arg in error message
- [IMPALA-2539](#) - Unmark collections slots of empty union operands
- [IMPALA-2554](#) - Change default buffer size for RPC servers and clients
- [IMPALA-2565](#) - Planner tests are flaky due to file size mismatches
- [IMPALA-2592](#) - DataStreamSender::Channel::CloseInternal() does not close the channel on an error
- [IMPALA-2599](#) - Pseudo-random sleep before acquiring kerberos ticket possibly not really pseudo-random
- [IMPALA-2711](#) - Fix memory leak in Rand()
- [IMPALA-2732](#) - Timestamp formats with non-padded values
- [IMPALA-2734](#) - Correlated EXISTS subqueries with HAVING clause return wrong results
- [IMPALA-2742](#) - Avoid unbounded MemPool growth with AcquireData()
- [IMPALA-2749](#) - Fix decimal multiplication overflow
- [IMPALA-2765](#) - Preserve return type of subexpressions substituted in isTrueWithNullSlots()
- [IMPALA-2788](#) - conv(bigint num, int from_base, int to_base) returns wrong result
- [IMPALA-2798](#) - Bring in AVRO-1617 fix and add test case for it
- [IMPALA-2818](#) - Fix cancellation crashes/hangs due to BlockOnWait() race
- [IMPALA-2820](#) - Support unquoted keywords as struct-field names
- [IMPALA-2832](#) - Fix cloning of FunctionCallExpr
- [IMPALA-2844](#) - Allow count(*) on RC files with complex types
- [IMPALA-2870](#) - Fix failing metadata.test_ddl.TestDdlStatements.test_create_table test
- [IMPALA-2894](#) - Move regression test into a different .test file
- [IMPALA-2906](#) - Fix an edge case with materializing TupleIsNullPredicates in analytic sorts
- [IMPALA-2914](#) - Fix DCHECK Check failed: HasDateOrTime()
- [IMPALA-2926](#) - Fix off-by-one bug in SelectNode::CopyRows()

- [IMPALA-2940](#) - Fix leak of dictionaries in Parquet scanner
- [IMPALA-3000](#) - Fix BitReader::Reset()
- [IMPALA-3034](#) - Verify all consumed memory of a MemTracker is always released at destruction time
- [IMPALA-3047](#) - Separate create table test with nested types
- [IMPALA-3054](#) - Disable probe side filters when spilling
- [IMPALA-3071](#) - Fix assignment of On-clause predicates belonging to an inner join
- [IMPALA-3085](#) - Unregister data sinks' MemTrackers at their Close() functions
- [IMPALA-3093](#) - ReopenClient() could NULL out 'client_key' causing a crash
- [IMPALA-3095](#) - Add configurable whitelist of authorized internal principals
- [IMPALA-3151](#) - Impala crash for avro table when casting to char data type
- [IMPALA-3194](#) - Allow queries materializing scalar type columns in RC/sequence files

Issues Fixed in Impala for CDH 5.6.0

The set of fixes for Impala in CDH 5.6.0 is the same as in CDH 5.5.2. See [Issues Fixed in Impala for CDH 5.5.2](#) on page 2427 for details.

Issues Fixed in Impala for CDH 5.5.6

For the full list of fixed issues for all CDH components, see [Issues Fixed in CDH 5.5.4](#) on page 2304.



Note: Impala 2.3.x is available as part of CDH 5.5.x and is not available for CDH 4. Cloudera does not intend to release future versions of Impala for CDH 4 outside patch and maintenance releases if required. Given the end-of-maintenance status for CDH 4, Cloudera recommends all customers to migrate to a recent CDH 5 release.

- [IMPALA-1928](#) - Fix Thrift client transport wrapping order
- [IMPALA-3369](#) - Add ALTER TABLE SET COLUMN STATS statement
- [IMPALA-3378](#) - HiveUdfCall::Open() produces unsynchronized access to JNIUtil::global_refs_vector
- [IMPALA-3379](#) - HBaseTableWriter::CreatePutList() produces unsynchronized access to JNIUtil::global_refs_vector
- [IMPALA-3441](#) - Check for malformed Avro data
- [IMPALA-3499](#) - Split catalog update
- [IMPALA-3575](#) - Add retry to backend connection request and rpc timeout
- [IMPALA-3633](#) - Cancel fragment if coordinator is gone
- [IMPALA-3682](#) - Do not retry unrecoverable socket creation errors
- [IMPALA-3687](#) - Prefer Avro field name during schema reconciliation
- [IMPALA-3698](#) - Fix Isilon permissions test
- [IMPALA-3711](#) - Remove unnecessary privilege checks in getDbsMetadata()
- [IMPALA-3732](#) - Handle string length overflow in Avro files
- [IMPALA-3751](#) - Fix clang build errors and warnings
- [IMPALA-3915](#) - Register privilege and audit requests when analyzing resolved table refs.
- [IMPALA-4135](#) - Thrift threaded server times-out connections during high load
- [IMPALA-4153](#) - Return valid non-NULL pointer for 0-byte allocations

Issues Fixed in Impala for CDH 5.5.4

For the full list of fixed issues for all CDH components, see [Issues Fixed in CDH 5.5.4](#) on page 2304.



Note: Impala 2.3.x is available as part of CDH 5.5.x and is not available for CDH 4. Cloudera does not intend to release future versions of Impala for CDH 4 outside patch and maintenance releases if required. Given the end-of-maintenance status for CDH 4, Cloudera recommends all customers to migrate to a recent CDH 5 release.

- [IMPALA-852](#) - ,IMPALA-2215: Analyze HAVING clause before aggregation

- [IMPALA-1092](#) - Fix estimates for trivial coord-only queries
- [IMPALA-1170](#) - Fix URL parsing when path contains '@'
- [IMPALA-1934](#) - Allow shell to retrieve LDAP password from shell cmd
- [IMPALA-2093](#) - Disallow NOT IN aggregate subqueries with a constant lhs expr
- [IMPALA-2184](#) - Don't inline timestamp methods with try/catch blocks in IR
- [IMPALA-2425](#) - Broadcast join hint not enforced when low memory limit is set
- [IMPALA-2503](#) - Add missing String.format() arg in error message
- [IMPALA-2539](#) - Unmark collections slots of empty union operands
- [IMPALA-2554](#) - Change default buffer size for RPC servers and clients
- [IMPALA-2565](#) - Planner tests are flaky due to file size mismatches
- [IMPALA-2592](#) - DataStreamSender::Channel::CloseInternal() does not close the channel on an error
- [IMPALA-2599](#) - Pseudo-random sleep before acquiring kerberos ticket possibly not really pseudo-random
- [IMPALA-2711](#) - Fix memory leak in Rand()
- [IMPALA-2719](#) - test_parquet_max_page_header fails on Isilon
- [IMPALA-2732](#) - Timestamp formats with non-padded values
- [IMPALA-2734](#) - Correlated EXISTS subqueries with HAVING clause return wrong results
- [IMPALA-2742](#) - Avoid unbounded MemPool growth with AcquireData()
- [IMPALA-2749](#) - Fix decimal multiplication overflow
- [IMPALA-2765](#) - Preserve return type of subexpressions substituted in isTrueWithNullSlots()
- [IMPALA-2788](#) - conv(bigint num, int from_base, int to_base) returns wrong result
- [IMPALA-2798](#) - Bring in AVRO-1617 fix and add test case for it
- [IMPALA-2818](#) - Fix cancellation crashes/hangs due to BlockOnWait() race
- [IMPALA-2820](#) - Support unquoted keywords as struct-field names
- [IMPALA-2832](#) - Fix cloning of FunctionCallExpr
- [IMPALA-2844](#) - Allow count(*) on RC files with complex types
- [IMPALA-2870](#) - Fix failing metadata.test_ddl.TestDdlStatements.test_create_table test
- [IMPALA-2894](#) - Move regression test into a different .test file
- [IMPALA-2906](#) - Fix an edge case with materializing TuplelsNullPredicates in analytic sorts
- [IMPALA-2914](#) - Fix DCHECK Check failed: HasDateOrTime()
- [IMPALA-2926](#) - Fix off-by-one bug in SelectNode::CopyRows()
- [IMPALA-2940](#) - Fix leak of dictionaries in Parquet scanner
- [IMPALA-3000](#) - Fix BitReader::Reset()
- [IMPALA-3034](#) - Verify all consumed memory of a MemTracker is always released at destruction time
- [IMPALA-3047](#) - Separate create table test with nested types
- [IMPALA-3054](#) - Disable probe side filters when spilling
- [IMPALA-3071](#) - Fix assignment of On-clause predicates belonging to an inner join
- [IMPALA-3085](#) - Unregister data sinks' MemTrackers at their Close() functions
- [IMPALA-3093](#) - ReopenClient() could NULL out 'client_key' causing a crash
- [IMPALA-3095](#) - Add configurable whitelist of authorized internal principals
- [IMPALA-3151](#) - Impala crash for avro table when casting to char data type
- [IMPALA-3194](#) - Allow queries materializing scalar type columns in RC/sequence files

Issues Fixed in Impala for CDH 5.5.2

This section lists the most serious or frequently encountered customer issues fixed in CDH 5.5.2 / Impala 2.3.2. For the full list of fixed Impala issues, see [Issues Fixed in CDH 5.5.2](#) on page 2310.

SEGV in AnalyticEvalNode touching NULL input_stream_

A query involving an analytic function could encounter a serious error. This issue was encountered infrequently, depending upon specific combinations of queries and data.

Bug: [IMPALA-2829](#)

Free local allocations per row batch in non-partitioned AGG and HJ

An outer join query could fail unexpectedly with an out-of-memory error when the “spill to disk” mechanism was turned off.

Bug: [IMPALA-2722](#)

Free local allocations once for every row batch when building hash tables

A join query could encounter a serious error due to an internal failure to allocate memory, which resulted in dereferencing a `NULL` pointer.

Bug: [IMPALA-2612](#)

Prevent migrating incorrectly inferred identity predicates into inline views

Referring to the same column twice in a view definition could cause the view to omit rows where that column contained a `NULL` value. This could cause incorrect results due to an inaccurate `COUNT (*)` value or rows missing from the result set.

Bug: [IMPALA-2643](#)

Fix GRANTS on URIs with uppercase letters

A `GRANT` statement for a URI could be ineffective if the URI contained uppercase letters, for example in an uppercase directory name. Subsequent statements, such as `CREATE EXTERNAL TABLE` with a `LOCATION` clause, could fail with an authorization exception.

Bug: [IMPALA-2695](#)

Avoid sending large partition stats objects over thrift

The `catalogd` daemon could encounter a serious error when loading the incremental statistics metadata for tables with large numbers of partitions and columns. The problem occurred when the internal representation of metadata for the table exceeded 2 GB, for example in a table with 20K partitions and 77 columns. The fix causes a `COMPUTE INCREMENTAL STATS` operation to fail if it would produce metadata that exceeded the maximum size.

Bug: [IMPALA-2664](#), [IMPALA-2648](#)

Throw `AnalysisError` if table properties are too large (for the Hive metastore)

`CREATE TABLE` or `ALTER TABLE` statements could fail with metastore database errors due to length limits on the `SERDEPROPERTIES` and `TBLPROPERTIES` clauses. (The limit on key size is 256, while the limit on value size is 4000.) The fix makes Impala handle these error conditions more cleanly, by detecting too-long values rather than passing them to the metastore database.

Bug: [IMPALA-2226](#)

Make `MAX_PAGE_HEADER_SIZE` configurable

Impala could fail to access Parquet data files with page headers larger than 8 MB, which could occur, for example, if the minimum or maximum values for a column were long strings. The fix adds a configuration setting `--max_page_header_size`, which you can use to increase the Impala size limit to a value higher than 8 MB.

Bug: [IMPALA-2273](#)

reduce scanner memory usage

Queries on Parquet tables could consume excessive memory (potentially multiple gigabytes) due to producing large intermediate data values while evaluating groups of rows. The workaround was to reduce the size of the `NUM_SCANNER_THREADS` query option, the `BATCH_SIZE` query option, or both.

Bug: [IMPALA-2473](#)

Handle error when distinct and aggregates are used with a having clause

A query that included a `DISTINCT` operator and a `HAVING` clause, but no aggregate functions or `GROUP BY`, would fail with an uninformative error message.

Bug: [IMPALA-2113](#)

Handle error when star based select item and aggregate are incorrectly used

A query that included * in the `SELECT` list, in addition to an aggregate function call, would fail with an uninformative message if the query had no `GROUP BY` clause.

Bug: [IMPALA-2225](#)

Refactor MemPool usage in HBase scan node

Queries involving HBase tables used substantially more memory than in earlier Impala versions. The problem occurred starting in Impala 2.2.8, as a result of the changes for IMPALA-2284. The fix for this issue involves removing a separate memory work area for HBase queries and reusing other memory that was already allocated.

Bug: [IMPALA-2731](#)

Fix migration/assignment of `ON`-clause predicates inside inline views

Some combinations of `ON` clauses in join queries could result in comparisons being applied at the wrong stage of query processing, leading to incorrect results. Wrong predicate assignment could happen under the following conditions:

- The query includes an inline view that contains an outer join.
- That inline view is joined with another table in the enclosing query block.
- That join has an `ON` clause containing a predicate that only references columns originating from the outer-joined tables inside the inline view.

Bug: [IMPALA-1459](#)

DCHECK in parquet scanner after block read error

A debug build of Impala could encounter a serious error after encountering some kinds of I/O errors for Parquet files. This issue only occurred in debug builds, not release builds.

Bug: [IMPALA-2558](#)

PAGG hits mem_limit when switching to I/O buffers

A join query could fail with an out-of-memory error despite the apparent presence of sufficient memory. The cause was the internal ordering of operations that could cause a later phase of the query to allocate memory required by an earlier phase of the query. The workaround was to either increase or decrease the `MEM_LIMIT` query option, because the issue would only occur for a specific combination of memory limit and data volume.

Bug: [IMPALA-2535](#)

Fix check failed: `sorter_runs_.back()->is_pinned_`

A query could fail with an internal error while calculating the memory limit. This was an infrequent condition uncovered during stress testing.

Bug: [IMPALA-2559](#)

Don't ignore Status returned by `DataStreamRecvr::CreateMerger()`

A query could fail with an internal error while calculating the memory limit. This was an infrequent condition uncovered during stress testing.

Bug: [IMPALA-2614](#), [IMPALA-2559](#)

`DataStreamSender::Send()` does not return an error status if `SendBatch()` failed

Bug: [IMPALA-2591](#)

Re-enable SSL and Kerberos on server-server

These fixes lift the restriction on using SSL encryption and Kerberos authentication together for internal communication between Impala components.

Bug: [IMPALA-2598](#), [IMPALA-2747](#)

Issues Fixed in Impala for CDH 5.5.1

The version of Impala that is included with CDH 5.5.1 / Impala 2.3.1 is identical to CDH 5.5.0 / Impala 2.3.0. There are no new bug fixes, new features, or incompatible changes.

Issues Fixed in Impala for CDH 5.5.0

This section lists the most serious or frequently encountered customer issues fixed in CDH 5.5.0 / Impala 2.3.0. Any issues already fixed in CDH 5.4 maintenance releases (up through CDH 5.4.8) are also included. Those issues are listed under the respective CDH 5.4 sections and are not repeated here. For the full list of fixed Impala issues, see [Issues Fixed in CDH 5.5.0](#) on page 2315.

Fixes for Serious Errors

A number of issues were resolved that could result in serious errors when encountered. The most critical or commonly encountered are listed here.

Bugs: [IMPALA-2168](#), [IMPALA-2378](#), [IMPALA-2369](#), [IMPALA-2357](#), [IMPALA-2319](#), [IMPALA-2314](#), [IMPALA-2016](#)

Fixes for Correctness Errors

A number of issues were resolved that could result in wrong results when encountered. The most critical or commonly encountered are listed here.

Bugs: [IMPALA-2192](#), [IMPALA-2440](#), [IMPALA-2090](#), [IMPALA-2086](#), [IMPALA-1947](#), [IMPALA-1917](#)

Issues Fixed in Impala for CDH 5.4.10

For the full list of fixed issues for all CDH components, see [Issues Fixed in CDH 5.4.10](#) on page 2323.



Note: The Impala 2.2.x maintenance releases now use the CDH 5.4.x numbering system rather than increasing the Impala version numbers. Impala 2.2 and higher are not available under CDH 4.

- [IMPALA-1702](#) - Check for duplicate table IDs at the end of analysis (issue not entirely fixed, but now fails gracefully)
- [IMPALA-2264](#) - Implicit casts to integers from decimals with higher precision sometimes allowed
- [IMPALA-2473](#) - Excessive memory usage by scan nodes
- [IMPALA-2621](#) - Fix flaky UNIX_TIMESTAMP() test
- [IMPALA-2643](#) - Nested inline view produces incorrect result when referencing the same column implicitly
- [IMPALA-2765](#) - AnalysisException: operands of type BOOLEAN and TIMESTAMP are not comparable when OUTER JOIN with CASE statement
- [IMPALA-2798](#) - After adding a column to avro table, Impala returns weird result if codegen is enabled.
- [IMPALA-2861](#) - Fix flaky scanner test added via IMPALA-2473 backport
- [IMPALA-2914](#) - Hit DCHECK Check failed: HasDateOrTime()
- [IMPALA-3034](#) - MemTracker leak on PHJ failure to spill
- [IMPALA-3085](#) - DataSinks' MemTrackers need to unregister themselves from parent
- [IMPALA-3093](#) - ReopenClient() could NULL out 'client_key' causing a crash
- [IMPALA-3095](#) - Allow additional Kerberos users to be authorized to access internal APIs

Issues Fixed in Impala for CDH 5.4.9

This section lists the most frequently encountered customer issues fixed in Impala for CDH 5.4.9.



Note: The Impala 2.2.x maintenance releases now use the CDH 5.4.x numbering system rather than increasing the Impala version numbers. Impala 2.2 and higher are not available under CDH 4.

For the full list of fixed issues, see [Issues Fixed in CDH 5.4.9](#) on page 2326.

Query return empty result if it contains NullLiteral in inlineview

If an inline view in a FROM clause contained a NULL literal, the result set was empty.

Bug: [IMPALA-1917](#)

HBase scan node uses 2-4x memory after upgrade to Impala 2.2.8

Queries involving HBase tables used substantially more memory than in earlier Impala versions. The problem occurred starting in Impala 2.2.8, as a result of the changes for IMPALA-2284. The fix for this issue involves removing a separate memory work area for HBase queries and reusing other memory that was already allocated.

Bug: [IMPALA-2731](#)

Fix migration/assignment of ON-clause predicates inside inline views

Some combinations of ON clauses in join queries could result in comparisons being applied at the wrong stage of query processing, leading to incorrect results. Wrong predicate assignment could happen under the following conditions:

- The query includes an inline view that contains an outer join.
- That inline view is joined with another table in the enclosing query block.
- That join has an ON clause containing a predicate that only references columns originating from the outer-joined tables inside the inline view.

Bug: [IMPALA-1459](#)

Fix wrong predicate assignment in outer joins

The join predicate for an OUTER JOIN clause could be applied at the wrong stage of query processing, leading to incorrect results.

Bug: [IMPALA-2446](#)

Avoid sending large partition stats objects over thrift

The `catalogd` daemon could encounter a serious error when loading the incremental statistics metadata for tables with large numbers of partitions and columns. The problem occurred when the internal representation of metadata for the table exceeded 2 GB, for example in a table with 20K partitions and 77 columns. The fix causes a `COMPUTE INCREMENTAL STATS` operation to fail if it would produce metadata that exceeded the maximum size.

Bug: [IMPALA-2648](#), [IMPALA-2664](#)

Avoid overflow when adding large intervals to TIMESTAMPS

Adding or subtracting a large INTERVAL value to a TIMESTAMP value could produce an incorrect result, with the value wrapping instead of returning an out-of-range error.

Bug: [IMPALA-1675](#)

Analysis exception when a binary operator contains an IN operator with values

An IN operator with literal values could cause a statement to fail if used as the argument to a binary operator, such as an equality test for a BOOLEAN value.

Bug: [IMPALA-1949](#)

Make MAX_PAGE_HEADER_SIZE configurable

Impala could fail to access Parquet data files with page headers larger than 8 MB, which could occur, for example, if the minimum or maximum values for a column were long strings. The fix adds a configuration setting `--max_page_header_size`, which you can use to increase the Impala size limit to a value higher than 8 MB.

Bug: [IMPALA-2273](#)

Fix spilling sorts with var-len slots that are NULL or empty.

A query that activated the spill-to-disk mechanism could fail if it contained a sort expression involving certain combinations of fixed-length or variable-length types.

Bug: [IMPALA-2357](#)

Work-around IMPALA-2344: Fail query with OOM in case block->Pin() fails

Some queries that activated the spill-to-disk mechanism could produce a serious error if there was insufficient memory to set up internal work areas. Now those queries produce normal out-of-memory errors instead.

Bug: [IMPALA-2344](#)

Crash (likely race) tearing down BufferedBlockMgr on query failure

A serious error could occur under rare circumstances, due to a race condition while freeing memory during heavily concurrent workloads.

Bug: [IMPALA-2252](#)

QueryExecState doesn't check for query cancellation or errors

A call to `setError()` in a user-defined function (UDF) would not cause the query to fail as expected.

Bug: [IMPALA-1746](#)

Impala throws `IllegalStateException` when inserting data into a partition while select subquery group by partition columns

An `INSERT ... SELECT` operation into a partitioned table could fail if the `SELECT` query included a `GROUP BY` clause referring to the partition key columns.

Bug: [IMPALA-2533](#)

Issues Fixed in Impala for CDH 5.4.8

This section lists the most frequently encountered customer issues fixed in Impala for CDH 5.4.8.



Note: The Impala 2.2.x maintenance releases now use the CDH 5.4.x numbering system rather than increasing the Impala version numbers. Impala 2.2 and higher are not available under CDH 4.

For the full list of fixed issues, see [Issues Fixed in CDH 5.4.8](#) on page 2328.

Impala is unable to read hive tables created with the "STORED AS AVRO" clause

Impala could not read Avro tables created in Hive with the `STORED AS AVRO` clause.

Bug: [IMPALA-1136](#), [IMPALA-2161](#)

make Parquet scanner fail query if the file size metadata is stale

If a Parquet file in HDFS was overwritten by a smaller file, Impala could encounter a serious error. Issuing a `INVALIDATE METADATA` statement before a subsequent query would avoid the error. The fix allows Impala to handle such inconsistencies in Parquet file length cleanly regardless of whether the table metadata is up-to-date.

Bug: [IMPALA-2213](#)

Avoid allocating `StringBuffer > 1GB` in `ScannerContext::Stream::GetBytesInternal()`

Impala could encounter a serious error when reading compressed text files larger than 1 GB. The fix causes Impala to issue an error message instead in this case.

Bug: [IMPALA-2249](#)

Disallow long ($1 << 30$) strings in `group_concat()`

A query using the `group_concat()` function could encounter a serious error if the returned string value was larger than 1 GB. Now the query fails with an error message in this case.

Bug: [IMPALA-2284](#)

avoid `FnvHash64to32` with empty inputs

An edge case in the algorithm used to distribute data among nodes could result in uneven distribution of work for some queries, with all data sent to the same node.

Bug: [IMPALA-2270](#)

The catalog does not close the connection to HMS during table invalidation

A communication error could occur between Impala and the Hive metastore database, causing Impala operations that update table metadata to fail.

Bug: [IMPALA-2348](#)

Wrong DCHECK in PHJ::ProcessProbeBatch

Certain queries could encounter a serious error if the spill-to-disk mechanism was activated.

Bug: [IMPALA-2364](#)

Avoid cardinality 0 in scan nodes of small tables and low selectivity

Impala could generate a suboptimal query plan for some queries involving small tables.

Bug: [IMPALA-2165](#)

Issues Fixed in Impala for CDH 5.4.7

This section lists the most frequently encountered customer issues fixed in Impala for CDH 5.4.7.



Note: The Impala 2.2.x maintenance releases now use the CDH 5.4.x numbering system rather than increasing the Impala version numbers. Impala 2.2 and higher are not available under CDH 4.

For the full list of fixed issues, see [Issues Fixed in CDH 5.4.7](#) on page 2331.

Warn if table stats are potentially corrupt.

Impala warns if it detects a discrepancy in table statistics: a table considered to have zero rows even though there are data files present. In this case, Impala also skips query optimizations that are normally applied to very small tables.

Bug: [IMPALA-1983:](#)

Pass correct child node in 2nd phase merge aggregation.

A query could encounter a serious error if it included a particular combination of aggregate functions and inline views.

Bug: [IMPALA-2266](#)

Set the output smap of an EmptySetNode produced from an empty inline view.

A query could encounter a serious error if it included an inline view whose subquery had no FROM clause.

Bug: [IMPALA-2216](#)

Set an InsertStmt's result exprs from the source statement's result exprs.

A CREATE TABLE AS SELECT or INSERT ... SELECT statement could produce different results than a SELECT statement, for queries including a FULL JOIN clause and including literal values in the select list.

Bug: [IMPALA-2203](#)

Fix planning of empty union operands with analytics.

A query could return incorrect results if it contained a UNION clause, calls to analytic functions, and a constant expression that evaluated to FALSE.

Bug: [IMPALA-2088](#)

Retain eq predicates bound by grouping slots with complex grouping exprs.

A query containing an INNER JOIN clause could return undesired rows. Some predicate specified in the ON clause could be omitted from the filtering operation.

Bug: [IMPALA-2089](#)

Row count not set for empty partition when spec is used with compute incremental stats

A `COMPUTE INCREMENTAL STATS` statement could leave the row count for an empty partition as -1, rather than initializing the row count to 0. The missing statistic value could result in reduced query performance.

Bug: [IMPALA-2199](#)

Explicit aliases + ordinals analysis bug

A query could encounter a serious error if it included column aliases with the same names as table columns, and used ordinal numbers in an `ORDER BY` or `GROUP BY` clause.

Bug: [IMPALA-1898](#)

Fix `TupleIsNotNullPredicate` to return false if no tuples are nullable.

A query could return incorrect results if it included an outer join clause, inline views, and calls to functions such as `coalesce()` that can generate `NULL` values.

Bug: [IMPALA-1987](#)

fix `Expr::ComputeResultsLayout()` logic

A query could return incorrect results if the table contained multiple `CHAR` columns with length of 2 or less, and the query included a `GROUP BY` clause that referred to multiple such columns.

Bug: [IMPALA-2178](#)

Substitute an `InsertStmt`'s partition key exprs with the root node's `smap`.

An `INSERT` statement could encounter a serious error if the `SELECT` portion called an analytic function.

Bug: [IMPALA-1737](#)

Issues Fixed in Impala for CDH 5.4.5

This section lists the most frequently encountered customer issues fixed in Impala for CDH 5.4.5.



Note: The Impala 2.2.x maintenance releases now use the CDH 5.4.x numbering system rather than increasing the Impala version numbers. Impala 2.2 and higher are not available under CDH 4.

For the full list of fixed issues, see [Issues Fixed in CDH 5.4.5](#) on page 2333.

Impala DML/DDL operations corrupt table metadata leading to Hive query failures

When the Impala `COMPUTE STATS` statement was run on a partitioned Parquet table that was created in Hive, the table subsequently became inaccessible in Hive. The table was still accessible to Impala. Regaining access in Hive required a workaround of creating a new table. The error displayed in Hive was:

```
Error: Error while compiling statement: FAILED: SemanticException Class not found:
com.cloudera.impala.hive.serde.ParquetInputFormat (state=42000,code=40000)
```

Bug: [IMPALA-2048](#)

Avoiding a `DCHECK` of `NULL` hash table in spilled right joins

A query could encounter a serious error if it contained a `RIGHT OUTER`, `RIGHT ANTI`, or `FULL OUTER` join clause and approached the memory limit on a host so that the “spill to disk” mechanism was activated.

Bug: [IMPALA-1929](#)

Bug in `PrintTColumnValue` caused wrong stats for `TINYINT` partition cols

Declaring a partition key column as a `TINYINT` caused problems with the `COMPUTE STATS` statement. The associated partitions would always have zero estimated rows, leading to potential inefficient query plans.

Bug: [IMPALA-2136](#)

Where clause does not propagate to joins inside nested views

A query that referred to a view whose query referred to another view containing a join, could return incorrect results. `WHERE` clauses for the outermost query were not always applied, causing the result set to include additional rows that should have been filtered out.

Bug: [IMPALA-2018](#)

Add `effective_user()` builtin

The `user()` function returned the name of the logged-in user, which might not be the same as the user name being checked for authorization if, for example, delegation was enabled.

Bug: [IMPALA-2064](#)

Resolution: Rather than change the behavior of the `user()` function, the fix introduces an additional function `effective_user()` that returns the user name that is checked during authorization.

Make UTC to local `TimestampValue` conversion faster.

Query performance was improved substantially for Parquet files containing `TIMESTAMP` data written by Hive, when the `-convert_legacy_hive_parquet_utc_timestamps=true` setting is in effect.

Bug: [IMPALA-2125](#)

Workaround [IMPALA-1619](#) in `BufferedBlockMgr::ConsumeMemory()`

A join query could encounter a serious error if the query approached the memory limit on a host so that the “spill to disk” mechanism was activated, and data volume in the join was large enough that an internal memory buffer exceeded 1 GB in size on a particular host. (Exceeding this limit would only happen for huge join queries, because Impala could split this intermediate data into 16 parts during the join query, and the buffer only contains compact bookkeeping data rather than the actual join column data.)

Bug: [IMPALA-2065](#)

Issues Fixed in Impala for CDH 5.4.3

This section lists the most frequently encountered customer issues fixed in Impala for CDH 5.4.3.



Note: The Impala 2.2.x maintenance releases now use the CDH 5.4.x numbering system rather than increasing the Impala version numbers. Impala 2.2 and higher are not available under CDH 4.

For the full list of fixed issues, see [Issues Fixed in CDH 5.4.3](#) on page 2337.

Enable using Isilon as the underlying filesystem.

The certification of CDH and Impala with the Isilon filesystem involves a number of fixes to performance and flexibility for dealing with I/O using remote reads. See [Using Impala with Isilon Storage](#) for details on using Impala and Isilon together.

Bug: [IMPALA-1968](#), [IMPALA-1730](#)

Expand set of supported timezones.

The set of timezones recognized by Impala was expanded. You can always find the latest list of supported timezones in the Impala source code, in the file [timezone_db.cc](#).

Bug: [IMPALA-1381](#)

Impala Timestamp ISO-8601 Support.

Impala can now process `TIMESTAMP` literals including a trailing `z`, signifying “Zulu” time, a synonym for UTC.

Bug: [IMPALA-1963](#)

Fix wrong warning when insert overwrite to empty table

An `INSERT OVERWRITE` operation would encounter an error if the `SELECT` portion of the statement returned zero rows, such as with a `LIMIT 0` clause.

Bug: [IMPALA-2008](#)

Expand parsing of decimals to include scientific notation

`DECIMAL` literals can now include `e` scientific notation. For example, now `CAST(1e3 AS DECIMAL(5,3))` is a valid expression. Formerly it returned `NULL`. Some scientific expressions might have worked before in `DECIMAL` context, but only when the scale was 0.

Bug: [IMPALA-1952](#)

Issues Fixed in Impala for CDH 5.4.1

This section lists the most frequently encountered customer issues fixed in Impala for CDH 5.4.1.



Note: The Impala 2.2.x maintenance releases now use the CDH 5.4.x numbering system rather than increasing the Impala version numbers. Impala 2.2 and higher are not available under CDH 4.

For the full list of fixed issues, see [Issues Fixed in CDH 5.4.1](#) on page 2339.

Issues Fixed in CDH 5.4 / Impala 2.2

This section lists the most frequently encountered customer issues fixed in Impala 2.2.0.

For the full list of fixed issues in Impala 2.2.0, including over 40 critical issues, see [this report in the JIRA system](#).



Note: Impala 2.2.0 is available as part of CDH 5.4.0 and is not available for CDH 4. Cloudera does not intend to release future versions of Impala for CDH 4 outside patch and maintenance releases if required. Given the end-of-maintenance status for CDH 4, Cloudera recommends all customers to migrate to a recent CDH 5 release.

Altering a column's type causes column stats to stop sticking for that column

When the type of a column was changed in either Hive or Impala through `ALTER TABLE CHANGE COLUMN`, the metastore database did not correctly propagate that change to the table that contains the column statistics. The statistics (particularly the `NDV`) for that column were permanently reset and could not be changed by Impala's `COMPUTE STATS` command. The underlying cause is a Hive bug (HIVE-9866).

Bug: [IMPALA-1607](#)

Resolution: Resolved by incorporating the fix for [HIVE-9866](#).

Workaround: On systems without the corresponding Hive fix, change the column back to its original type. The stats reappear and you can recompute or drop them.

Impala may leak or use too many file descriptors

If a file was truncated in HDFS without a corresponding `REFRESH` in Impala, Impala could allocate memory for file descriptors and not free that memory.

Bug: [IMPALA-1854](#)

Spurious stale block locality messages

Impala could issue messages stating the block locality metadata was stale, when the metadata was actually fine. The internal “remote bytes read” counter was not being reset properly. This issue did not cause an actual slowdown in query execution, but the spurious error could result in unnecessary debugging work and unnecessary use of the `INVALIDATE METADATA` statement.

Bug: [IMPALA-1712](#)

DROP TABLE fails after COMPUTE STATS and ALTER TABLE RENAME to a different database.

When a table was moved from one database to another, the column statistics were not pointed to the new database. This could result in lower performance for queries due to unavailable statistics, and also an inability to drop the table.

Bug: [IMPALA-1711](#)

IMPALA-1556 causes memory leak with secure connections

`impalad` daemons could experience a memory leak on clusters using Kerberos authentication, with memory usage growing as more data is transferred across the secure channel, either to the client program or between Impala nodes. The same issue affected LDAP-secured clusters to a lesser degree, because the LDAP security only covers data transferred back to client programs.

Bug: [IMPALA-1674](#)

`unix_timestamp()` does not return correct time

The `unix_timestamp()` function could return an incorrect value (a constant value of 1).

Bug: [IMPALA-1623](#)

Impala incorrectly handles text data missing a newline on the last line

Some queries did not recognize the final line of a text data file if the line did not end with a newline character. This could lead to inconsistent results, such as a different number of rows for `SELECT COUNT(*)` as opposed to `SELECT *`.

Bug: [IMPALA-1476](#)

Impala's ACLs check do not consider all group ACLs, only checked first one.

If the HDFS user ID associated with the `impalad` process had read or write access in HDFS based on group membership, Impala statements could still fail with HDFS permission errors if that group was not the first listed group for that user ID.

Bug: [IMPALA-1805](#)

Fix infinite loop opening or closing file with invalid metadata

Truncating a file in HDFS, after Impala had cached the file metadata, could produce a hang when Impala queried a table containing that file.

Bug: [IMPALA-1794](#)

Cannot write Parquet files when values are larger than 64KB

Impala could sometimes fail to `INSERT` into a Parquet table if a column value such as a `STRING` was larger than 64 KB.

Bug: [IMPALA-1705](#)

Impala Will Not Run on Certain Intel CPUs

This fix relaxes the CPU requirement for Impala. Now only the SSE3 instruction set is required. Formerly, SSE4.1 instructions were generated, making Impala refuse to start on some older CPUs.

Bug: [IMPALA-1646](#)

Issues Fixed in Impala for CDH 5.3.10

For the full list of fixed issues for all CDH components, see [Issues Fixed in CDH 5.3.10](#) on page 2346.



Note: This Impala maintenance release is only available as part of CDH 5, not under CDH 4.

- [IMPALA-1702](#) - "invalidate metadata" can cause duplicate TableIds (issue not entirely fixed, but now fails gracefully)
- [IMPALA-2125](#) - Improve perf when reading timestamps from parquet files written by hive

- [IMPALA-2565](#) - Planner tests are flaky due to file size mismatches
- [IMPALA-3095](#) - Allow additional Kerberos users to be authorized to access internal APIs

Issues Fixed in the 2.1.7 Release / CDH 5.3.9

This section lists the most significant Impala issues fixed in Impala 2.1.7 for CDH 5.3.9.

For the full list of Impala fixed issues in this release, see [Issues Fixed in CDH 5.3.9](#) on page 2348.



Note: This Impala maintenance release is only available as part of CDH 5, not under CDH 4.

Query return empty result if it contains NullLiteral in inlineview

If an inline view in a FROM clause contained a NULL literal, the result set was empty.

Bug: [IMPALA-1917](#)

Fix edge cases for decimal/integer cast

A value of type `DECIMAL(3,0)` could be incorrectly cast to `TINYINT`. The resulting out-of-range value could be incorrect. After the fix, the smallest type that is allowed for this cast is `INT`, and attempting to use `DECIMAL(3,0)` in a `TINYINT` context produces a “loss of precision” error.

Bug: [IMPALA-2264](#)

Constant filter expressions are not checked for errors and state cleanup on exception / DCHECK on destroying an ExprContext

An invalid constant expression in a WHERE clause (for example, an invalid regular expression pattern) could fail to clean up internal state after raising a query error. Therefore, certain combinations of invalid expressions in a query could cause a crash, or cause a query to continue when it should halt with an error.

Bug: [IMPALA-1756](#), [IMPALA-2514](#)

QueryExecState does not check for query cancellation or errors

A call to `setError()` in a user-defined function (UDF) would not cause the query to fail as expected.

Bug: [IMPALA-1746](#), [IMPALA-2141](#)

Issues Fixed in the 2.1.6 Release / CDH 5.3.8

This section lists the most significant Impala issues fixed in Impala 2.1.6 for CDH 5.3.8.

For the full list of Impala fixed issues in this release, see [Issues Fixed in CDH 5.3.8](#) on page 2350.



Note: This Impala maintenance release is only available as part of CDH 5, not under CDH 4.

Wrong DCHECK in PHJ::ProcessProbeBatch

Certain queries could encounter a serious error if the spill-to-disk mechanism was activated.

Bug: [IMPALA-2364](#)

LargestSpilledPartition was not checking if partition is closed

Certain queries could encounter a serious error if the spill-to-disk mechanism was activated.

Bug: [IMPALA-2314](#)

Avoid cardinality 0 in scan nodes of small tables and low selectivity

Impala could generate a suboptimal query plan for some queries involving small tables.

Bug: [IMPALA-2165](#)

fix Expr::ComputeResultsLayout() logic

Queries using the `GROUP BY` operator on multiple `CHAR` columns with length less than or equal to 2 characters could return incorrect results for some columns.

Bug: [IMPALA-2178](#)

Properly unescape string value for HBase filters

Queries against HBase tables could return incomplete results if the `WHERE` clause included string comparisons using literals containing escaped quotation marks.

Bug: [IMPALA-2133](#)

Avoiding a DCHECK of NULL hash table in spilled right joins

A query could encounter a serious error if it contained a `RIGHT OUTER`, `RIGHT ANTI`, or `FULL OUTER` join clause and approached the memory limit on a host so that the “spill to disk” mechanism was activated.

Bug: [IMPALA-1929](#)

Issues Fixed in the 2.1.5 Release / CDH 5.3.6

This section lists the most significant Impala issues fixed in Impala 2.1.5 for CDH 5.3.6.

For the full list of Impala fixed issues in this release, see [Issues Fixed in CDH 5.3.6](#) on page 2352.



Note: This Impala maintenance release is only available as part of CDH 5, not under CDH 4.

Avoid calling `ProcessBatch` with `out_batch->AtCapacity` in right joins

Queries including `RIGHT OUTER JOIN`, `RIGHT ANTI JOIN`, or `FULL OUTER JOIN` clauses and involving a high volume of data could encounter a serious error.

Bug: [IMPALA-1919](#)

Issues Fixed in the 2.1.4 Release / CDH 5.3.4

This section lists the most significant Impala issues fixed in Impala 2.1.4 for CDH 5.3.4. Because CDH 5.3.5 does not include any code changes for Impala, Impala 2.1.4 is included with both CDH 5.3.4 and 5.3.5.

For the full list of Impala fixed issues in Impala 2.1.4 for CDH 5.3.4, see [Issues Fixed in CDH 5.3.4](#) on page 2354.



Note: This Impala maintenance release is only available as part of CDH 5, not under CDH 4.

Crash: `impala::TuplesNullPredicate::Prepare`

When expressions that tested for `NULL` were used in combination with analytic functions, an error could occur. (The original crash issue was fixed by an earlier patch.)

Bug: [IMPALA-1519](#)

Expand parsing of decimals to include scientific notation

`DECIMAL` literals could include `e` scientific notation. For example, now `CAST(1e3 AS DECIMAL(5,3))` is a valid expression. Formerly it returned `NULL`. Some scientific expressions might have worked before in `DECIMAL` context, but only when the scale was 0.

Bug: [IMPALA-1952](#)

INSERT/CTAS evaluates and applies constant predicates.

An `INSERT OVERWRITE` statement would write new data, even if a constant clause such as `WHERE 1 = 0` should have prevented it from writing any rows.

Bug: [IMPALA-1860](#)

Assign predicates below analytic functions with a compatible partition by clause

If the `PARTITION BY` clause in an analytic function refers to partition key columns in a partitioned table, now Impala can perform partition pruning during the analytic query.

Bug: [IMPALA-1900](#)

FIRST_VALUE may produce incorrect results with preceding windows

A query using the `FIRST_VALUE` analytic function and a window defined with the `PRECEDING` keyword could produce wrong results.

Bug: [IMPALA-1888](#)

FIRST_VALUE rewrite fn type might not match slot type

A query referencing a `DECIMAL` column with the `FIRST_VALUE` analytic function could encounter an error.

Bug: [IMPALA-1559](#)

AnalyticEvalNode cannot handle partition/order by exprs with NaN

A query using an analytic function could encounter an error if the evaluation of an analytic `ORDER BY` or `PARTITION` expression resulted in a NaN value, for example if the `ORDER BY` or `PARTITION` contained a division operation where both operands were zero.

Bug: [IMPALA-1808](#)

AnalyticEvalNode not properly handling nullable tuples

An analytic function containing only an `OVER` clause could encounter an error if another part of the query (specifically an outer join) produced all-NULL tuples.

Bug: [IMPALA-1562](#)

Issues Fixed in the 2.1.3 Release / CDH 5.3.3

This section lists the most significant issues fixed in Impala 2.1.3.

For the full list of fixed issues in Impala 2.1.3, see [Issues Fixed in CDH 5.3.3](#) on page 2356.



Note: Impala 2.1.3 is available as part of CDH 5.3.3, not under CDH 4.

Add compatibility flag for Hive-Parquet-Timestamps

When Hive writes `TIMESTAMP` values, it represents them in the local time zone of the server. Impala expects `TIMESTAMP` values to always be in the UTC time zone, possibly leading to inconsistent results depending on which component created the data files. This patch introduces a new startup flag, `-convert_legacy_hive_parquet_utc_timestamps` for the `impalad` daemon. Specify `-convert_legacy_hive_parquet_utc_timestamps=true` to make Impala recognize Parquet data files written by Hive and automatically adjust `TIMESTAMP` values read from those files into the UTC time zone for compatibility with other Impala `TIMESTAMP` processing. Although this setting is currently turned off by default, consider enabling it if practical in your environment, for maximum interoperability with Hive-created Parquet files.

Bug: [IMPALA-1658](#)

Use `snprintf()` instead of `lexical_cast()` in float-to-string casts

Converting a floating-point value to a `STRING` could be slower than necessary.

Bug: [IMPALA-1738](#)

Fix partition spilling cleanup when new stream OOMs

Certain calls to aggregate functions with `STRING` arguments could encounter a serious error when the system ran low on memory and attempted to activate the spill-to-disk mechanism. The error message referenced the function `impala::AggregateFunctions::StringValGetValue`.

Bug: [IMPALA-1865](#)

Impala's ACLs check do not consider all group ACLs, only checked first one.

If the HDFS user ID associated with the `impalad` process had read or write access in HDFS based on group membership, Impala statements could still fail with HDFS permission errors if that group was not the first listed group for that user ID.

Bug: [IMPALA-1805](#)

Fix infinite loop opening or closing file with invalid metadata

Truncating a file in HDFS, after Impala had cached the file metadata, could produce a hang when Impala queried a table containing that file.

Bug: [IMPALA-1794](#)

external-data-source-executor leaking global jni refs

Successive calls to the data source API could result in excessive memory consumption, with memory allocated but never freed.

Bug: [IMPALA-1801](#)

Spurious stale block locality messages

Impala could issue messages stating the block locality metadata was stale, when the metadata was actually fine. The internal "remote bytes read" counter was not being reset properly. This issue did not cause an actual slowdown in query execution, but the spurious error could result in unnecessary debugging work and unnecessary use of the `INVALIDATE METADATA` statement.

Bug: [IMPALA-1712](#)

Issues Fixed in the 2.1.2 Release / CDH 5.3.2

This section lists the most significant issues fixed in Impala 2.1.2.

For the full list of fixed issues in Impala 2.1.2, see [this report in the JIRA system](#).



Note: Impala 2.1.2 is available as part of CDH 5.3.2, not under CDH 4.

Impala incorrectly handles double numbers with more than 19 significant decimal digits

When a floating-point value was read from a text file and interpreted as a `FLOAT` or `DOUBLE` value, it could be incorrectly interpreted if it included more than 19 significant digits.

Bug: [IMPALA-1622](#)

`unix_timestamp()` does not return correct time

The `unix_timestamp()` function could return an incorrect value (a constant value of 1).

Bug: [IMPALA-1623](#)

Row Count Mismatch: Partition pruning with `NULL`

A query against a partitioned table could return incorrect results if the `WHERE` clause compared the partition key to `NULL` using operators such as `=` or `!=`.

Bug: [IMPALA-1535](#)

Fetch column stats in bulk using new (Hive .13) HMS APIs

The performance of the `COMPUTE STATS` statement and queries was improved, particularly for wide tables.

Bug: [IMPALA-1120](#)

Issues Fixed in the 2.1.1 Release / CDH 5.3.1

This section lists the most significant issues fixed in Impala 2.1.1.

For the full list of fixed issues in Impala 2.1.1, see [this report in the JIRA system](#).

[IMPALA-1556 causes memory leak with secure connections](#)

`impalad` daemons could experience a memory leak on clusters using Kerberos authentication, with memory usage growing as more data is transferred across the secure channel, either to the client program or between Impala nodes. The same issue affected LDAP-secured clusters to a lesser degree, because the LDAP security only covers data transferred back to client programs.

Bug: [IMPALA-1674](#)

[TSaslServerTransport::Factory::getTransport\(\) leaks transport map entries](#)

`impalad` daemons in clusters secured by Kerberos or LDAP could experience a slight memory leak on each connection. The accumulation of unreleased memory could cause problems on long-running clusters.

Bug: [IMPALA-1668](#)

Issues Fixed in the 2.1.0 Release / CDH 5.3.0

This section lists the most significant issues fixed in Impala 2.1.0.

For the full list of fixed issues in Impala 2.1.0, see [this report in the JIRA system](#).

[Kerberos fetches 3x slower](#)

Transferring large result sets back to the client application on Kerberos

Bug: [IMPALA-1455](#)

[Compressed file needs to be hold on entirely in Memory](#)

Queries on gzipped text files required holding the entire data file and its uncompressed representation in memory at the same time. `SELECT` and `COMPUTE STATS` statements could fail or perform inefficiently as a result. The fix enables streaming reads for gzipped text, so that the data is uncompressed as it is read.

Bug: [IMPALA-1556](#)

[Cannot read hbase metadata with NullPointerException: null](#)

Impala might not be able to access HBase tables, depending on the associated levels of Impala and HBase on the system.

Bug: [IMPALA-1611](#)

[Serious errors / crashes](#)

Improved code coverage in Impala testing uncovered a number of potentially serious errors that could occur with specific query syntax. These errors are resolved in Impala 2.1.

Bug: [IMPALA-1553](#), [IMPALA-1528](#), [IMPALA-1526](#), [IMPALA-1524](#), [IMPALA-1508](#), [IMPALA-1493](#), [IMPALA-1501](#), [IMPALA-1483](#)

Issues Fixed in the 2.0.5 Release / CDH 5.2.6

For the full list of fixed issues in Impala 2.0.5, see [this report in the JIRA system](#).



Note: Impala 2.0.5 is available as part of CDH 5.2.6, not under CDH 4.

Issues Fixed in the 2.0.4 Release / CDH 5.2.5

This section lists the most significant issues fixed in Impala 2.0.4.

For the full list of fixed issues in Impala 2.0.4, see [this report in the JIRA system](#).



Note: Impala 2.0.4 is available as part of CDH 5.2.5, not under CDH 4.

Add compatibility flag for Hive-Parquet-Timestamps

When Hive writes `TIMESTAMP` values, it represents them in the local time zone of the server. Impala expects `TIMESTAMP` values to always be in the UTC time zone, possibly leading to inconsistent results depending on which component created the data files. This patch introduces a new startup flag, `-convert_legacy_hive_parquet_utc_timestamps` for the `impalad` daemon. Specify `-convert_legacy_hive_parquet_utc_timestamps=true` to make Impala recognize Parquet data files written by Hive and automatically adjust `TIMESTAMP` values read from those files into the UTC time zone for compatibility with other Impala `TIMESTAMP` processing. Although this setting is currently turned off by default, consider enabling it if practical in your environment, for maximum interoperability with Hive-created Parquet files.

Bug: [IMPALA-1658](#)

IoMgr infinite loop opening/closing file when shorter than cached metadata size

If a table data file was replaced by a shorter file outside of Impala, such as with `INSERT OVERWRITE` in Hive producing an empty output file, subsequent Impala queries could hang.

Bug: [IMPALA-1794](#)

Issues Fixed in the 2.0.3 Release / CDH 5.2.4

This section lists the most significant issues fixed in Impala 2.0.3.

For the full list of fixed issues in Impala 2.0.3, see [this report in the JIRA system](#).



Note: Impala 2.0.3 is available as part of CDH 5.2.4, not under CDH 4.

Anti join could produce incorrect results when spilling

An anti-join query (or a `NOT EXISTS` operation that was rewritten internally into an anti-join) could produce incorrect results if Impala reached its memory limit, causing the query to write temporary results to disk.

Bug: [IMPALA-1471](#)

Row Count Mismatch: Partition pruning with NULL

A query against a partitioned table could return incorrect results if the `WHERE` clause compared the partition key to `NULL` using operators such as `=` or `!=`.

Bug: [IMPALA-1535](#)

Fetch column stats in bulk using new (Hive .13) HMS APIs

The performance of the `COMPUTE STATS` statement and queries was improved, particularly for wide tables.

Bug: [IMPALA-1120](#)

Issues Fixed in the 2.0.2 Release / CDH 5.2.3

This section lists the most significant issues fixed in Impala 2.0.2.

For the full list of fixed issues in Impala 2.0.2, see [this report in the JIRA system](#).



Note: Impala 2.0.2 is available as part of CDH 5.2.3, not under CDH 4.

GROUP BY on STRING column produces inconsistent results

Some operations in queries submitted through Hue or other HiveServer2 clients could produce inconsistent results.

Bug: [IMPALA-1453](#)

Fix leaked file descriptor and excessive file descriptor use

Impala could encounter an error from running out of file descriptors. The fix reduces the amount of time file descriptors are kept open, and avoids leaking file descriptors when read operations encounter errors.

unix_timestamp() does not return correct time

The `unix_timestamp()` function could return a constant value 1 instead of a representation of the time.

Bug: [IMPALA-1623](#)

Impala should randomly select cached replica

To avoid putting too heavy a load on any one node, Impala now randomizes which scan node processes each HDFS data block rather than choosing the first cached block replica.

Bug: [IMPALA-1586](#)

Impala does not always give short name to Llama.

In clusters secured by Kerberos or LDAP, a discrepancy in internal transmission of user names could cause a communication error with Llama.

Bug: [IMPALA-1606](#)

accept unmangled native UDF symbols

The `CREATE FUNCTION` statement could report that it could not find a function entry point within the `.so` file for a UDF written in C++, even if the corresponding function was present.

Bug: [IMPALA-1475](#)

Issues Fixed in the 2.0.1 Release / CDH 5.2.1

This section lists the most significant issues fixed in Impala 2.0.1.

For the full list of fixed issues in Impala 2.0.1, see [this report in the JIRA system](#).

Queries fail with metastore exception after upgrade and compute stats

After running the `COMPUTE STATS` statement on an Impala table, subsequent queries on that table could fail with the exception message `Failed to load metadata for table: default.stats_test`.

Bug: [IMPALA-1416](#)

Workaround: Upgrading to CDH 5.2.1, or another level of CDH that includes the fix for HIVE-8627, prevents the problem from affecting future `COMPUTE STATS` statements. On affected levels of CDH, or for Impala tables that have become inaccessible, the workaround is to disable the `hive.metastore.try.direct.sql` setting in the Hive metastore `hive-site.xml` file and issue the `INVALIDATE METADATA` statement for the affected table. You do not need to rerun the `COMPUTE STATS` statement for the table.

Issues Fixed in the 2.0.0 Release / CDH 5.2.0

This section lists the most significant issues fixed in Impala 2.0.0.

For the full list of fixed issues in Impala 2.0.0, see [this report in the JIRA system](#).

Join Hint is dropped when used inside a view

Hints specified within a view query did not take effect when the view was queried, leading to slow performance. As part of this fix, Impala now supports hints embedded within comments.

Bug: [IMPALA-995](#)"

WHERE condition ignored in simple query with RIGHT JOIN

Potential wrong results for some types of queries.

Bug: [IMPALA-1101](#)"

Query with self joined table may produce incorrect results

Potential wrong results for some types of queries.

Bug: [IMPALA-1102](#)"

Incorrect plan after reordering predicates (inner join following outer join)

Potential wrong results for some types of queries.

Bug: [IMPALA-1118](#)"

Combining fragments with compatible data partitions can lead to incorrect results due to type incompatibilities (missing casts).

Potential wrong results for some types of queries.

Bug: [IMPALA-1123](#)"

Predicate dropped: Inline view + DISTINCT aggregate in outer query

Potential wrong results for some types of queries.

Bug: [IMPALA-1165](#)"

Reuse of a column in JOIN predicate may lead to incorrect results

Potential wrong results for some types of queries.

Bug: [IMPALA-1353](#)"

Usage of TRUNC with string timestamp reliably crashes node

Serious error for certain combinations of function calls and data types.

Bug: [IMPALA-1105](#)"

Timestamp Cast Returns invalid TIMESTAMP

Serious error for certain combinations of function calls and data types.

Bug: [IMPALA-1109](#)"

IllegalStateException upon JOIN of DECIMAL columns with different precision

DECIMAL columns with different precision could not be compared in join predicates.

Bug: [IMPALA-1121](#)"

Allow creating Avro tables without column definitions. Allow `COMPUTE STATS` to always work on Impala-created Avro tables.

Hive-created Avro tables with columns specified by a JSON file or literal could produce errors when queried in Impala, and could not be used with the `COMPUTE STATS` statement. Now you can create such tables in Impala to avoid such errors.

Bug: [IMPALA-1104](#)"

Ensure all webserver output is escaped

The Impala debug web UI did not properly encode all output.

Bug: [IMPALA-1133](#)"

Queries with union in inline view have empty resource requests

Certain queries could run without obeying the limits imposed by resource management.

Bug: [IMPALA-1236](#)"

Impala does not employ ACLs when checking path permissions for `LOAD` and `INSERT`

Certain `INSERT` and `LOAD DATA` statements could fail unnecessarily, if the target directories in HDFS had restrictive HDFS permissions, but those permissions were overridden by HDFS extended ACLs.

Bug: [IMPALA-1279](#)"

Impala does not map principals to lowercase, affecting Sentry authorisation

In a Kerberos environment, the principal name was not mapped to lowercase, causing issues when a user logged in with an uppercase principal name and Sentry authorization was enabled.

Bug: [IMPALA-1334](#)"

Issues Fixed in the 1.4.4 Release / CDH 5.1.5

For the list of fixed issues, see [Issues Fixed in CDH 5.1.5](#) in the *CDH 5 Release Notes*.



Note: Impala 1.4.4 is available as part of CDH 5.1.5, not under CDH 4.

Issues Fixed in the 1.4.3 Release / CDH 5.1.4

Impala 1.4.3 includes fixes to address what is known as the POODLE vulnerability in SSLv3. SSLv3 access is disabled in the Impala debug web UI.



Note: Impala 1.4.3 is available as part of CDH 5.1.4, and under CDH 4.

Issues Fixed in the 1.4.2 Release / CDH 5.1.3

This section lists the most significant issues fixed in Impala 1.4.2.

For the full list of fixed issues in Impala 1.4.2, see [this report in the JIRA system](#).



Note: Impala 1.4.3 is available as part of CDH 5.1.4, and under CDH 4.

Issues Fixed in the 1.4.1 Release / CDH 5.1.2

This section lists the most significant issues fixed in Impala 1.4.1.

For the full list of fixed issues in Impala 1.4.1, see [this report in the JIRA system](#).



Note: Impala 1.4.1 is only available as part of CDH 5.1.2, not under CDH 4.

impalad terminating with Boost exception

Occasionally, a non-trivial query run through Llama could encounter a serious error. The detailed error in the log was:

```
boost::exception_detail::clone_impl
<boost::exception_detail::error_info_injector<boost::lock_error> >
```

Severity: High

Impalad uses wrong string format when writing logs

Impala log files could contain internal error messages due to a problem formatting certain strings. The messages consisted of a Java call stack starting with:

```
jni-util.cc:177] java.util.MissingFormatArgumentException: Format specifier 's'
```

Update HS2 client API.

A downlevel version of the HiveServer2 API could cause difficulty retrieving the precision and scale of a DECIMAL value.

Bug: [IMPALA-1107](#)

Impalad catalog updates can fail with error: "IllegalArgumentException: fromKey out of range" at com.cloudera.impala.catalog.CatalogDeltaLog

The error in the title could occur following a DDL statement. This issue was discovered during internal testing and has not been reported in customer environments.

Bug: [IMPALA-1093](#)

"Total" time counter does not capture all the network transmit time

The time for some network operations was not counted in the report of total time for a query, making it difficult to diagnose network-related performance issues.

Bug: [IMPALA-1131](#)

Impala will crash when reading certain Avro files containing bytes data

Certain Avro fields for byte data could cause Impala to be unable to read an Avro data file, even if the field was not part of the Impala table definition. With this fix, Impala can now read these Avro data files, although Impala queries cannot refer to the "bytes" fields.

Bug: [IMPALA-1149](#)

Support specifying a custom AuthorizationProvider in Impala

The `--authorization_policy_provider_class` option for `impalad` was added back. This option specifies a custom `AuthorizationProvider` class rather than the default `HadoopGroupAuthorizationProvider`. It had been used for internal testing, then removed in Impala 1.4.0, but it was considered useful by some customers.

Bug: [IMPALA-1142](#)

Issues Fixed in the 1.4.0 Release / CDH 5.1.0

This section lists the most significant issues fixed in Impala 1.4.0.

For the full list of fixed issues in Impala 1.4.0, see [this report in the JIRA system](#).

Failed DCHECK in disk-io-mgr-reader-context.cc:174

The serious error in the title could occur, with the supplemental message:

```
num_used_buffers_ < 0: #used=-1 during cancellation HDFS cached data
```

The issue was due to the use of HDFS caching with data files accessed by Impala. Support for HDFS caching in Impala was introduced in Impala 1.4.0 for CDH 5.1.0. The fix for this issue was backported to Impala 1.3.x, and is the only change in Impala 1.3.2 for CDH 5.0.4.

Bug: [IMPALA-1019](#)

Workaround: On CDH 5.0.x, upgrade to CDH 5.0.4 with Impala 1.3.2, where this issue is fixed. In Impala 1.3.0 or 1.3.1 on CDH 5.0.x, do not use HDFS caching for Impala data files in Impala internal or external tables. If some of these data files are cached (for example because they are used by other components that take advantage of HDFS caching), set the query option `DISABLE_CACHED_READS=true`. To set that option for all Impala queries across all sessions, start `impalad` with the `-default_query_options` option and include this setting in the option argument, or on a cluster managed by Cloudera Manager, fill in this option setting on the **Impala Daemon** options page.

Resolution: This issue is fixed in Impala 1.3.2 for CDH 5.0.4. The addition of HDFS caching support in Impala 1.4 means that this issue does not apply to any new level of Impala on CDH 5.

impala-shell only works with ASCII characters

The `impala-shell` interpreter could encounter errors processing SQL statements containing non-ASCII characters.

Bug: [IMPALA-489](#)

The extended view definition SQL text in Views created by Impala should always have fully-qualified table names

When a view was accessed while inside a different database, references to tables were not resolved unless the names were fully qualified when the view was created.

Bug: [IMPALA-962](#)

Impala forgets about partitions with non-existent locations

If an `ALTER TABLE` specified a non-existent HDFS location for a partition, afterwards Impala would not be able to access the partition at all.

Bug: [IMPALA-741](#)

CREATE TABLE LIKE fails if source is a view

The `CREATE TABLE LIKE` clause was enhanced to be able to create a table with the same column definitions as a view. The resulting table is a text table unless the `STORED AS` clause is specified, because a view does not have an associated file format to inherit.

Bug: [IMPALA-834](#)

Improve partition pruning time

Operations on tables with many partitions could be slow due to the time to evaluate which partitions were affected. The partition pruning code was speeded up substantially.

Bug: [IMPALA-887](#)

Improve compute stats performance

The performance of the `COMPUTE STATS` statement was improved substantially. The efficiency of its internal operations was improved, and some statistics are no longer gathered because they are not currently used for planning Impala queries.

Bug: [IMPALA-1003](#)

When I run `CREATE TABLE new_table LIKE avro_table`, the schema does not get mapped properly from an avro schema to a hive schema

After a `CREATE TABLE LIKE` statement using an Avro table as the source, the new table could have incorrect metadata and be inaccessible, depending on how the original Avro table was created.

Bug: [IMPALA-185](#)

Race condition in IoMgr. Blocked ranges enqueued after cancel.

Impala could encounter a serious error after a query was cancelled.

Bug: [IMPALA-1046](#)

Deadlock in scan node

A deadlock condition could make all `impalad` daemons hang, making the cluster unresponsive for Impala queries.

Bug: [IMPALA-1083](#)

Issues Fixed in the 1.3.3 Release / CDH 5.0.5

Impala 1.3.3 includes fixes to address what is known as the POODLE vulnerability in SSLv3. SSLv3 access is disabled in the Impala debug web UI.



Note: Impala 1.3.3 is only available as part of CDH 5.0.5, not under CDH 4.

Issues Fixed in the 1.3.2 Release / CDH 5.0.4

This backported bug fix is the only change between Impala 1.3.1 and Impala 1.3.2.



Note: Impala 1.3.3 is only available as part of CDH 5.0.5, not under CDH 4.

Failed DCHECK in disk-io-mgr-reader-context.cc:174

The serious error in the title could occur, with the supplemental message:

```
num_used_buffers_ < 0: #used=-1 during cancellation HDFS cached data
```

The issue was due to the use of HDFS caching with data files accessed by Impala. Support for HDFS caching in Impala was introduced in Impala 1.4.0 for CDH 5.1.0. The fix for this issue was backported to Impala 1.3.x, and is the only change in Impala 1.3.2 for CDH 5.0.4.

Bug: [IMPALA-1019](#)

Workaround: On CDH 5.0.x, upgrade to CDH 5.0.4 with Impala 1.3.2, where this issue is fixed. In Impala 1.3.0 or 1.3.1 on CDH 5.0.x, do not use HDFS caching for Impala data files in Impala internal or external tables. If some of these data files are cached (for example because they are used by other components that take advantage of HDFS caching), set the query option `DISABLE_CACHED_READS=true`. To set that option for all Impala queries across all sessions, start `impalad` with the `-default_query_options` option and include this setting in the option argument, or on a cluster managed by Cloudera Manager, fill in this option setting on the **Impala Daemon** options page.

Resolution: This issue is fixed in Impala 1.3.2 for CDH 5.0.4. The addition of HDFS caching support in Impala 1.4 means that this issue does not apply to any new level of Impala on CDH 5.

Issues Fixed in the 1.3.1 Release / CDH 5.0.3

This section lists the most significant issues fixed in Impala 1.3.1.

For the full list of fixed issues in Impala 1.3.1, see [this report in the JIRA system](#).

Impalad crashes when left joining inline view that has aggregate using distinct

Impala could encounter a severe error in a query combining a left outer join with an inline view containing a `COUNT(DISTINCT)` operation.

Bug: [IMPALA-904](#)

Incorrect result with group by query with null value in group by data

If the result of a `GROUP BY` operation is `NULL`, the resulting row might be omitted from the result set. This issue depends on the data values and data types in the table.

Bug: [IMPALA-901](#)

Drop Function does not clear local library cache

When a UDF is dropped through the `DROP FUNCTION` statement, and then the UDF is re-created with a new `.so` library or JAR file, the original version of the UDF is still used when the UDF is called from queries.

Bug: [IMPALA-786](#)

Workaround: Restart the `impalad` daemon on all nodes.

Compute stats doesn't propagate underlying error correctly

If a `COMPUTE STATS` statement encountered an error, the error message is "Query aborted" with no further detail. Common reasons why a `COMPUTE STATS` statement might fail include network errors causing the coordinator node to lose contact with other `impalad` instances, and column names that match Impala [reserved words](#). (Currently, if a column name is an Impala reserved word, `COMPUTE STATS` always returns an error.)

Bug: [IMPALA-762](#)

Inserts should respect changes in partition location

After an `ALTER TABLE` statement that changes the `LOCATION` property of a partition, a subsequent `INSERT` statement would always use a path derived from the base data directory for the table.

Bug: [IMPALA-624](#)

Text data with carriage returns generates wrong results for `count(*)`

A `COUNT(*)` operation could return the wrong result for text tables using nul characters (ASCII value 0) as delimiters.

Bug: [IMPALA-13](#)

Workaround: Impala adds support for ASCII 0 characters as delimiters through the clause `FIELDS TERMINATED BY '\0'`.

IO Mgr should take instance memory limit into account when creating io buffers

Impala could allocate more memory than necessary during certain operations.

Bug: [IMPALA-488](#)

Workaround: Before issuing a `COMPUTE STATS` statement for a Parquet table, reduce the number of threads used in that operation by issuing `SET NUM_SCANNER_THREADS=2` in `impala-shell`. Then issue `UNSET NUM_SCANNER_THREADS` before continuing with queries.

Impala should provide an option for new sub directories to automatically inherit the permissions of the parent directory

When new subdirectories are created underneath a partitioned table by an `INSERT` statement, previously the new subdirectories always used the default HDFS permissions for the `impala` user, which might not be suitable for directories intended to be read and written by other components also.

Bug: [IMPALA-827](#)

Resolution: In Impala 1.3.1 and higher, you can specify the `--insert_inherit_permissions` configuration when starting the `impalad` daemon.

Illegal state exception (or crash) in query with UNION in inline view

Impala could encounter a severe error in a query where the `FROM` list contains an inline view that includes a `UNION`. The exact type of the error varies.

Bug: [IMPALA-888](#)

INSERT column reordering doesn't work with SELECT clause

The ability to specify a subset of columns in an `INSERT` statement, with order different than in the target table, was not working as intended.

Bug: [IMPALA-945](#)

Issues Fixed in the 1.3.0 Release / CDH 5.0.0

This section lists the most significant issues fixed in Impala 1.3.0, primarily issues that could cause wrong results, or cause problems running the `COMPUTE STATS` statement, which is very important for performance and scalability.

For the full list of fixed issues, see [this report in the JIRA system](#).

Inner join after right join may produce wrong results

The automatic join reordering optimization could incorrectly reorder queries with an outer join or semi join followed by an inner join, producing incorrect results.

Bug: [IMPALA-860](#)

Workaround: Including the `STRAIGHT_JOIN` keyword in the query prevented the issue from occurring.

Incorrect results with codegen on multi-column group by with NULLs.

A query with a `GROUP BY` clause referencing multiple columns could introduce incorrect `NULL` values in some columns of the result set. The incorrect `NULL` values could appear in rows where a different `GROUP BY` column actually did return `NULL`.

Bug: [IMPALA-850](#)

Using distinct inside aggregate function may cause incorrect result when using having clause

A query could return incorrect results if it combined an aggregate function call, a `DISTINCT` operator, and a `HAVING` clause, without a `GROUP BY` clause.

Bug: [IMPALA-845](#)

Aggregation on union inside (inline) view not distributed properly.

An aggregation query or a query with `ORDER BY` and `LIMIT` could be executed on a single node in some cases, rather than distributed across the cluster. This issue affected queries whose `FROM` clause referenced an inline view containing a `UNION`.

Bug: [IMPALA-831](#)

Wrong expression may be used in aggregate query if there are multiple similar expressions

If a `GROUP BY` query referenced the same columns multiple times using different operators, result rows could contain multiple copies of the same expression.

Bug: [IMPALA-817](#)

Incorrect results when changing the order of aggregates in the select list with codegen enabled

Referencing the same columns in both a `COUNT()` and a `SUM()` call in the same query, or some other combinations of aggregate function calls, could incorrectly return a result of 0 from one of the aggregate functions. This issue affected references to `TINYINT` and `SMALLINT` columns, but not `INT` or `BIGINT` columns.

Bug: [IMPALA-765](#)

Workaround: Setting the query option `DISABLE_CODEGEN=TRUE` prevented the incorrect results. Switching the order of the function calls could also prevent the issue from occurring.

Union queries give Wrong result in a UNION followed by SIGSEGV in another union

A UNION query could produce a wrong result, followed by a serious error for a subsequent UNION query.

Bug: [IMPALA-723](#)

String data in MR-produced parquet files may be read incorrectly

Impala could return incorrect string results when reading uncompressed Parquet data files containing multiple row groups. This issue only affected Parquet data files produced by MapReduce jobs.

Bug: [IMPALA-729](#)

Compute stats need to use quotes with identifiers that are Impala keywords

Using a column or table name that conflicted with Impala keywords could prevent running the COMPUTE STATS statement for the table.

Bug: [IMPALA-777](#)

COMPUTE STATS child queries do not inherit parent query options.

The COMPUTE STATS statement did not use the setting of the MEM_LIMIT query option in impala-shell, potentially causing problems gathering statistics for wide Parquet tables.

Bug: [IMPALA-903](#)

COMPUTE STATS should update partitions in batches

The COMPUTE STATS statement could be slow or encounter a timeout while analyzing a table with many partitions.

Bug: [IMPALA-880](#)

Fail early (in analysis) when COMPUTE STATS is run against Avro table with no columns

If the columns for an Avro table were all defined in the TBLPROPERTIES or SERDEPROPERTIES clauses, the COMPUTE STATS statement would fail after completely analyzing the table, potentially causing a long delay. Although the COMPUTE STATS statement still does not work for such tables, now the problem is detected and reported immediately.

Bug: [IMPALA-867](#)

Workaround: Re-create the Avro table with columns defined in SQL style, using the output of SHOW CREATE TABLE. (See the JIRA page for detailed steps.)

Issues Fixed in the 1.2.4 Release

This section lists the most significant issues fixed in Impala 1.2.4. For the full list of fixed issues, see [this report in the JIRA system](#).

The Catalog Server exits with an OOM error after a certain number of CREATE statements

A large number of concurrent CREATE TABLE statements can cause the catalogd process to consume excessive memory, and potentially be killed due to an out-of-memory condition.

Bug: [IMPALA-818](#)

Workaround: Restart the catalogd service and re-try the DDL operations that failed.

Catalog Server consumes excessive cpu cycle

A large number of tables and partitions could result in unnecessary CPU overhead during Impala idle time and background operations.

Bug: [IMPALA-821](#)

Resolution: Catalog server processing was optimized in several ways.

Query against Avro table crashes Impala with codegen enabled

A query against a TIMESTAMP column in an Avro table could encounter a serious issue.

Bug: [IMPALA-828](#)**Workaround:** Set the query option `DISABLE_CODEGEN=TRUE`

Statestore seems to send concurrent heartbeats to the same subscriber leading to repeated "Subscriber '*hostname*' is registering with statestore, ignoring update" messages

Impala nodes could produce repeated error messages after recovering from a communication error with the statestore service.

Bug: [IMPALA-809](#)

Join predicate incorrectly ignored

A join query could produce wrong results if multiple equality comparisons between the same tables referred to the same column.

Bug: [IMPALA-805](#)

Query result differing between Impala and Hive

Certain outer join queries could return wrong results. If one of the tables involved in the join was an inline view, some tests from the `WHERE` clauses could be applied to the wrong phase of the query.

[ArrayIndexOutOfBoundsException / Invalid query handle when reading large HBase cell](#)

An HBase cell could contain a value larger than 32 KB, leading to a serious error when Impala queries that table. The error could occur even if the applicable row is not part of the result set.

Bug: [IMPALA-715](#)**Workaround:** Use smaller values in the HBase table, or exclude the column containing the large value from the result set.

[select with distinct and full outer join, impalad coredump](#)

A query involving a `DISTINCT` operator combined with a `FULL OUTER JOIN` could encounter a serious error.

Bug: [IMPALA-735](#)**Workaround:** Set the query option `DISABLE_CODEGEN=TRUE`

Impala cannot load tables with more than `Short.MAX_VALUE` number of partitions

If a table had more than 32,767 partitions, Impala would not recognize the partitions above the 32K limit and query results could be incomplete.

Bug: [IMPALA-749](#)

Various issues with HBase row key specification

Queries against HBase tables could fail with an error if the row key was compared to a function return value rather than a string constant. Also, queries against HBase tables could fail if the `WHERE` clause contained combinations of comparisons that could not possibly match any row key.

Resolution: Queries now return appropriate results when function calls are used in the row key comparison. For queries involving non-existent row keys, such as `WHERE row_key IS NULL` or where the lower bound is greater than the upper bound, the query succeeds and returns an empty result set.

Issues Fixed in the 1.2.3 Release

This release is a fix release that supercedes Impala 1.2.2, with the same features and fixes as 1.2.2 plus one additional fix for compatibility with Parquet files generated outside of Impala by components such as Hive, Pig, or MapReduce.

[Impala cannot read Parquet files with multiple row groups](#)

The `parquet-mr` library included with CDH4.5 writes files that are not readable by Impala, due to the presence of multiple row groups. Queries involving these data files might result in a crash or a failure with an error such as "Column chunk should not contain two dictionary pages".

This issue does not occur for Parquet files produced by Impala `INSERT` statements, because Impala only produces files with a single row group.

Bug: [IMPALA-720](#)

Issues Fixed in the 1.2.2 Release

This section lists the most significant issues fixed in Impala 1.2.2. For the full list of fixed issues, see [this report in the JIRA system](#).

Order of table references in FROM clause is critical for optimal performance

Impala does not currently optimize the join order of queries; instead, it joins tables in the order in which they are listed in the FROM clause. Queries that contain one or more large tables on the right hand side of joins (either an explicit join expressed as a JOIN statement or a join implicit in the list of table references in the FROM clause) may run slowly or crash Impala due to out-of-memory errors. For example:

```
SELECT ... FROM small_table JOIN large_table
```

Anticipated Resolution: Fixed in Impala 1.2.2.

Workaround: In Impala 1.2.2 and higher, use the `COMPUTE STATS` statement to gather statistics for each table involved in the join query, after data is loaded. Prior to Impala 1.2.2, modify the query, if possible, to join the largest table first. For example:

```
SELECT ... FROM small_table JOIN large_table
```

should be modified to:

```
SELECT ... FROM large_table JOIN small_table
```

Parquet in CDH4.5 writes data files that are sometimes unreadable by Impala

Some Parquet files could be generated by other components that Impala could not read.

Bug: [IMPALA-694](#)

Resolution: The underlying issue is being addressed by a fix in the CDH Parquet libraries. Impala 1.2.2 works around the problem and reads the existing data files.

Deadlock in statestore when unregistering a subscriber and building a topic update

The statestore service could experience an internal error leading to a hang.

Bug: [IMPALA-699](#)

IllegalStateException when doing a union involving a group by

A `UNION` query where one side involved a `GROUP BY` operation could cause a serious error.

Bug: [IMPALA-687](#)

Impala Parquet Writer hit DCHECK in RleEncoder

A serious error could occur when doing an `INSERT` into a Parquet table.

Bug: [IMPALA-689](#)

Hive UDF jars cannot be loaded by the FE

If the JAR file for a Java-based Hive UDF was not in the `CLASSPATH`, the UDF could not be called during a query.

Bug: [IMPALA-695](#)

Issues Fixed in the 1.2.1 Release

This section lists the most significant issues fixed in Impala 1.2.1. For the full list of fixed issues, see [this report in the JIRA system](#).

Scanners use too much memory when reading past scan range

While querying a table with long column values, Impala could over-allocate memory leading to an out-of-memory error. This problem was observed most frequently with tables using uncompressed RCFile or text data files.

Bug: [IMPALA-525](#)

Resolution: Fixed in 1.2.1

Join node consumes memory way beyond mem-limit

A join query could allocate a temporary work area that was larger than needed, leading to an out-of-memory error. The fix makes Impala return unused memory to the system when the memory limit is reached, avoiding unnecessary memory errors.

Bug: [IMPALA-657](#)

Resolution: Fixed in 1.2.1

Excessive memory consumption when query tables with 1k columns (Parquet file)

Impala could encounter an out-of-memory condition setting up work areas for Parquet tables with many columns. The fix reduces the size of the allocated memory when not actually needed to hold table data.

Bug: [IMPALA-652](#)

Resolution: Fixed in 1.2.1

Issues Fixed in the 1.2.0 Beta Release

This section lists the most significant issues fixed in Impala 1.2 (beta). For the full list of fixed issues, see [this report in the JIRA system](#).

Issues Fixed in the 1.1.1 Release

This section lists the most significant issues fixed in Impala 1.1.1. For the full list of fixed issues, see [this report in the JIRA system](#).

Unexpected LLVM Crash When Querying Doubles on CentOS 5.x

Certain queries involving `DOUBLE` columns could fail with a serious error. The fix improves the generation of native machine instructions for certain chipsets.

Bug: [IMPALA-477](#)

"block size is too big" error with Snappy-compressed RCFile containing null

Queries could fail with a "block size is too big" error, due to `NULL` values in RCFile tables using Snappy compression.

Bug: [IMPALA-482](#)

Cannot query RC file for table that has more columns than the data file

Queries could fail if an Impala RCFile table was defined with more columns than in the corresponding RCFile data files.

Bug: [IMPALA-510](#)

Views Sometimes Not Utilizing Partition Pruning

Certain combinations of clauses in a view definition for a partitioned table could result in inefficient performance and incorrect results.

Bug: [IMPALA-495](#)

Update the serde name we write into the metastore for Parquet tables

The SerDes class string written into Parquet data files created by Impala was updated for compatibility with Parquet support in Hive. See [Incompatible Changes Introduced in Impala 1.1.1](#) on page 1936 for the steps to update older Parquet data files for Hive compatibility.

Bug: [IMPALA-485](#)

Selective queries over large tables produce unnecessary memory consumption

A query returning a small result sets from a large table could tie up memory unnecessarily for the duration of the query.

Bug: [IMPALA-534](#)

Impala stopped to query AVRO tables

Queries against Avro tables could fail depending on whether the Avro schema URL was specified in the `TBLPROPERTIES` or `SERDEPROPERTIES` field. The fix causes Impala to check both fields for the schema URL.

Bug: [IMPALA-538](#)

Impala continues to allocate more memory even though it has exceed its mem-limit

Queries could allocate substantially more memory than specified in the `impalad -mem_limit` startup option. The fix causes more frequent checking of the limit during query execution.

Bug: [IMPALA-520](#)

Issues Fixed in the 1.1.0 Release

This section lists the most significant issues fixed in Impala 1.1. For the full list of fixed issues, see [this report in the JIRA system](#).

10-20% perf regression for most queries across all table formats

This issue is due to a performance tradeoff between systems running many queries concurrently, and systems running a single query. Systems running only a single query could experience lower performance than in early beta releases. Systems running many queries simultaneously should experience higher performance than in the beta releases.

planner fails with "Join requires at least one equality predicate between the two tables" when "from" table order does not match "where" join order

A query could fail if it involved 3 or more tables and the last join table was specified as a subquery.

Bug: [IMPALA-85](#)

Parquet writer uses excessive memory with partitions

`INSERT` statements against partitioned tables using the Parquet format could use excessive amounts of memory as the number of partitions grew large.

Bug: [IMPALA-257](#)

Comments in `impala-shell` in interactive mode are not handled properly causing syntax errors or wrong results

The `impala-shell` interpreter did not accept comment entered at the command line, making it problematic to copy and paste from scripts or other code examples.

Bug: [IMPALA-192](#)

Cancelled queries sometimes aren't removed from the inflight query list

The Impala web UI would sometimes display a query as if it were still running, after the query was cancelled.

Bug: [IMPALA-364](#)

Impala's 1.0.1 Shell Broke Python 2.4 Compatibility (AttributeError: 'module' object has no attribute 'field_size_limit')

The `impala-shell` command in Impala 1.0.1 does not work with Python 2.4, which is the default on Red Hat 5.

For the `impala-shell` command in Impala 1.0, the `-o` option (pipe output to a file) does not work with Python 2.4.

Bug: [IMPALA-396](#)

Issues Fixed in the 1.0.1 Release

This section lists the most significant issues fixed in Impala 1.0.1. For the full list of fixed issues, see [this report in the JIRA system](#).

Impala parquet scanner cannot read all data files generated by other frameworks

Impala might issue an erroneous error message when processing a Parquet data file produced by a non-Impala Hadoop component.

Bug: [IMPALA-333](#)

Resolution: Fixed

Impala is unable to query RCFile tables which describe fewer columns than the file's header.

If an RCFile table definition had fewer columns than the fields actually in the data files, queries would fail.

Bug: [IMPALA-293](#)

Resolution: Fixed

Impala does not correctly substitute `_HOST` with `hostname` in `--principal`

The `_HOST` placeholder in the `--principal` startup option was not substituted with the correct hostname, potentially leading to a startup error in setups using Kerberos authentication.

Bug: [IMPALA-351](#)

Resolution: Fixed

HBase query missed the last region

A query for an HBase table could omit data from the last region.

Bug: [IMPALA-356](#)

Resolution: Fixed

Hbase region changes are not handled correctly

After a region in an HBase table was split or moved, an Impala query might return incomplete or out-of-date results.

Bug: [IMPALA-300](#)

Resolution: Fixed

Query state for successful create table is EXCEPTION

After a successful `CREATE TABLE` statement, the corresponding query state would be incorrectly reported as `EXCEPTION`.

Bug: [IMPALA-349](#)

Resolution: Fixed

Double check release of JNI-allocated byte-strings

Operations involving calls to the Java JNI subsystem (for example, queries on HBase tables) could allocate memory but not release it.

Bug: [IMPALA-358](#)

Resolution: Fixed

Impala returns 0 for bad time values in `UNIX_TIMESTAMP`, Hive returns `NULL`

Impala returns 0 for bad time values in `UNIX_TIMESTAMP`, Hive returns `NULL`.

Impala:

```
impala> select UNIX_TIMESTAMP('10:02:01') ;
impala> 0
```

Hive:

```
hive> select UNIX_TIMESTAMP('10:02:01') FROM tmp;
hive> NULL
```

Bug: [IMPALA-16](#)**Anticipated Resolution:** Fixed

INSERT INTO TABLE SELECT <constant> does not work.

Insert INTO TABLE SELECT <constant> will not insert any data and may return an error.

Anticipated Resolution: Fixed

Issues Fixed in the 1.0 GA Release

Here are the major user-visible issues fixed in Impala 1.0. For a full list of fixed issues, see [this report in the public issue tracker](#).

Undeterministically receive "ERROR: unknown row batch destination..." and "ERROR: Invalid query handle" from impala shell when running union query

A query containing both UNION and LIMIT clauses could intermittently cause the `impalad` process to halt with a segmentation fault.

Bug: [IMPALA-183](#)**Resolution:** Fixed

Insert with NULL partition keys results in SIGSEGV.

An INSERT statement specifying a NULL value for one of the partitioning columns could cause the `impalad` process to halt with a segmentation fault.

Bug: [IMPALA-190](#)**Resolution:** Fixed

INSERT queries don't show completed profiles on the debug webpage

In the Impala web user interface, the profile page for an INSERT statement showed obsolete information for the statement once it was complete.

Bug: [IMPALA-217](#)**Resolution:** Fixed

Impala HBase scan is very slow

Queries involving an HBase table could be slower than expected, due to excessive memory usage on the Impala nodes.

Bug: [IMPALA-231](#)**Resolution:** FixedAdd some library version validation logic to `impalad` when loading `impala-lzo` shared library

No validation was done to check that the `impala-lzo` shared library was compatible with the version of Impala, possibly leading to a crash when using LZ0-compressed text files.

Bug: [IMPALA-234](#)**Resolution:** Fixed

Workaround: Always upgrade the `impala-lzo` library at the same time as you upgrade Impala itself.

Problems inserting into tables with `TIMESTAMP` partition columns leading table metadata loading failures and failed dchecks

`INSERT` statements for tables partitioned on columns involving datetime types could appear to succeed, but cause errors for subsequent queries on those tables. The problem was especially serious if an improperly formatted timestamp value was specified for the partition key.

Bug: [IMPALA-238](#)

Resolution: Fixed

Ctrl-C sometimes interrupts shell in system call, rather than cancelling query

Pressing Ctrl-C in the `impala-shell` interpreter could sometimes display an error and return control to the shell, making it impossible to cancel the query.

Bug: [IMPALA-243](#)

Resolution: Fixed

Empty string partition value causes metastore update failure

Specifying an empty string or `NULL` for a partition key in an `INSERT` statement would fail.

Bug: [IMPALA-252](#)

Resolution: Fixed. The behavior for empty partition keys was made more compatible with the corresponding Hive behavior.

`Round()` does not output the right precision

The `round()` function did not always return the correct number of significant digits.

Bug: [IMPALA-266](#)

Resolution: Fixed

Cannot cast string literal to string

Casting from a string literal back to the same type would cause an “invalid type cast” error rather than leaving the original value unchanged.

Bug: [IMPALA-267](#)

Resolution: Fixed

Excessive mem usage for certain queries which are very selective

Some queries that returned very few rows experienced unnecessary memory usage.

Bug: [IMPALA-288](#)

Resolution: Fixed

HdfsScanNode crashes in UpdateCounters

A serious error could occur for relatively small and inexpensive queries.

Bug: [IMPALA-289](#)

Resolution: Fixed

Parquet performance issues on large dataset

Certain aggregation queries against Parquet tables were inefficient due to lower than required thread utilization.

Bug: [IMPALA-292](#)

Resolution: Fixed

impala not populating hive metadata correctly for create table

The Impala `CREATE TABLE` command did not fill in the `owner` and `tbl_type` columns in the Hive metastore database.

Bug: [IMPALA-295](#)

Resolution: Fixed. The metadata was made more Hive-compatible.

impala daemons die if statestore goes down

The `impalad` instances in a cluster could halt when the `statestore` process became unavailable.

Bug: [IMPALA-312](#)

Resolution: Fixed

Constant `SELECT` clauses do not work in subqueries

A subquery would fail if the `SELECT` statement inside it returned a constant value rather than querying a table.

Bug: [IMPALA-67](#)

Resolution: Fixed

Right outer Join includes `NULL`s as well and hence wrong result count

The result set from a right outer join query could include erroneous rows containing `NULL` values.

Bug: [IMPALA-90](#)

Resolution: Fixed

Parquet scanner hangs for some queries

The Parquet scanner non-deterministically hangs when executing some queries.

Bug: [IMPALA-204](#)

Resolution: Fixed

Issues Fixed in Version 0.7 of the Beta Release

Impala does not gracefully handle unsupported Hive table types (`INDEX` and `VIEW` tables)

When attempting to load metadata from an unsupported Hive table type (`INDEX` and `VIEW` tables), Impala fails with an unclear error message.

Bug: [IMPALA-167](#)

Resolution: Fixed in 0.7

DDL statements (`CREATE/ALTER/DROP TABLE`) are not supported in the Impala Beta Release

Resolution: Fixed in 0.7

Avro is not supported in the Impala Beta Release

Resolution: Fixed in 0.7

Workaround: None

Impala does not currently allow limiting the memory consumption of a single query

It is currently not possible to limit the memory consumption of a single query. All tables on the right hand side of `JOIN` statements need to be able to fit in memory. If they do not, Impala may crash due to out of memory errors.

Resolution: Fixed in 0.7

Aggregate of a subquery result set returns wrong results if the subquery contains a 'limit' and data is distributed across multiple nodes

Aggregate of a subquery result set returns wrong results if the subquery contains a 'limit' clause and data is distributed across multiple nodes. From the query plan, it looks like we are just summing the results from each worker node.

Bug: [IMPALA-20](#)**Resolution:** Fixed in 0.7

Partition pruning for arbitrary predicates that are fully bound by a particular partition column

We currently cannot utilize a predicate like "country_code in ('DE', 'FR', 'US')" to do partitioning pruning, because that requires an equality predicate or a binary comparison.

We should create a superclass of `planner.ValueRange`, `ValueSet`, that can be constructed with an arbitrary predicate, and whose `isInRange(analyzer, valueExpr)` constructs a literal predicate by substitution of the `valueExpr` into the predicate.

Bug: [IMPALA-144](#)**Resolution:** Fixed in 0.7

Issues Fixed in Version 0.6 of the Beta Release

[Impala reads the NameNode address and port as command line parameters](#)

Impala reads the NameNode address and port as command line parameters rather than reading them from `core-site.xml`. Updating the NameNode address in the `core-site.xml` file does not propagate to Impala.

Severity: Low

Resolution: Fixed in 0.6 - Impala reads the namenode location and port from the Hadoop configuration files, though setting `-nn` and `-nn_port` overrides this. Users are advised not to set `-nn` or `-nn_port`.

[Queries may fail on secure environment due to impalad Kerberos ticket expiration](#)

Queries may fail on secure environment due to `impalad` Kerberos tickets expiring. This can happen if the Impala `-kerberos_reinit_interval` flag is set to a value ten minutes or less. This may lead to an `impalad` requesting a ticket with a lifetime that is less than the time to the next ticket renewal.

Bug: [IMPALA-64](#)**Resolution:** Fixed in 0.6[Concurrent queries may fail when Impala uses Thrift to communicate with the Hive Metastore](#)

Concurrent queries may fail when Impala is using Thrift to communicate with part of the Hive Metastore such as the Hive Metastore Service. In such a case, the error `get_fields failed: out of sequence response` may occur because Impala shared a single Hive Metastore Client connection across threads. With Impala 0.6, a separate connection is used for each metadata request.

Bug: [IMPALA-48](#)**Resolution:** Fixed in 0.6[impalad fails to start if unable to connect to the Hive Metastore](#)

Impala fails to start if it is unable to establish a connection with the Hive Metastore. This behavior was fixed, allowing Impala to start, even when no Metastore is available.

Bug: [IMPALA-58](#)**Resolution:** Fixed in 0.6[Impala treats database names as case-sensitive in some contexts](#)

In some queries (including "USE database" statements), database names are treated as case-sensitive. This may lead queries to fail with an `IllegalStateException`.

Bug: [IMPALA-44](#)**Resolution:** Fixed in 0.6

Impala does not ignore hidden HDFS files

Impala does not ignore hidden HDFS files, meaning those files prefixed with a period '.' or underscore '_'. This diverges from Hive/MapReduce, which skips these files.

Bug: [IMPALA-18](#)

Resolution: Fixed in 0.6

Issues Fixed in Version 0.5 of the Beta Release

Impala may have reduced performance on tables that contain a large number of partitions

Impala may have reduced performance on tables that contain a large number of partitions. This is due to extra overhead reading/parsing the partition metadata.

Resolution: Fixed in 0.5

Backend client connections not getting cached causes an observable latency in secure clusters

Backend impalads do not cache connections to the coordinator. On a secure cluster, this introduces a latency proportional to the number of backend clients involved in query execution, as the cost of establishing a secure connection is much higher than in the non-secure case.

Bug: [IMPALA-38](#)

Resolution: Fixed in 0.5

Concurrent queries may fail with error: "Table object has not been initialised : `PARTITIONS`"

Concurrent queries may fail with error: "Table object has not been initialised : `PARTITIONS`". This was due to a lack of locking in the Impala table/database metadata cache.

Bug: [IMPALA-30](#)

Resolution: Fixed in 0.5

UNIX_TIMESTAMP format behaviour deviates from Hive when format matches a prefix of the time value

The Impala UNIX_TIMESTAMP(val, format) operation compares the length of format and val and returns NULL if they do not match. Hive instead effectively truncates val to the length of the format parameter.

Bug: [IMPALA-15](#)

Resolution: Fixed in 0.5

Issues Fixed in Version 0.4 of the Beta Release

Impala fails to refresh the Hive metastore if a Hive temporary configuration file is removed

Impala is impacted by Hive bug [HIVE-3596](#) which may cause metastore refreshes to fail if a Hive temporary configuration file is deleted (normally located at /tmp/hive-<user>-<tmp_number>.xml). Additionally, the impala-shell will incorrectly report that the failed metadata refresh completed successfully.

Anticipated Resolution: To be fixed in a future release

Workaround: Restart the `impalad` service. Use the `impalad` log to check for metadata refresh errors.

lpad/rpad builtin functions is not correct.

The `lpad/rpad` builtin functions generate the wrong results.

Resolution: Fixed in 0.4

Files with .gz extension reported as 'not supported'

Compressed files with extensions incorrectly generate an exception.

Bug: [IMPALA-14](#)

Resolution: Fixed in 0.4

Queries with large limits would hang.

Some queries with large limits were hanging.

Resolution: Fixed in 0.4

Order by on a string column produces incorrect results if there are empty strings

Resolution: Fixed in 0.4

Issues Fixed in Version 0.3 of the Beta Release

All table loading errors show as unknown table

If Impala is unable to load the metadata for a table for any reason, a subsequent query referring to that table will return an `unknown table` error message, even if the table is known.

Resolution: Fixed in 0.3

A table that cannot be loaded will disappear from SHOW TABLES

After failing to load metadata for a table, Impala removes that table from the list of known tables returned in `SHOW TABLES`. Subsequent attempts to query the table returns 'unknown table', even if the metadata for that table is fixed.

Resolution: Fixed in 0.3

Impala cannot read from HBase tables that are not created as external tables in the hive metastore.

Attempting to select from these tables fails.

Resolution: Fixed in 0.3

Certain queries that contain OUTER JOINs may return incorrect results

Queries that contain OUTER JOINs may not return the correct results if there are predicates referencing any of the joined tables in the WHERE clause.

Resolution: Fixed in 0.3.

Issues Fixed in Version 0.2 of the Beta Release

Subqueries which contain aggregates cannot be joined with other tables or Impala may crash

Subqueries that contain an aggregate cannot be joined with another table or Impala may crash. For example:

```
SELECT * FROM (SELECT sum(col1) FROM some_table GROUP BY col1) t1 JOIN other_table ON (...);
```

Resolution: Fixed in 0.2

An insert with a limit that runs as more than one query fragment inserts more rows than the limit.

For example:

```
INSERT OVERWRITE TABLE test SELECT * FROM test2 LIMIT 1;
```

Resolution: Fixed in 0.2

Query with limit clause might fail.

For example:

```
SELECT * FROM test2 LIMIT 1;
```

Resolution: Fixed in 0.2

Files in unsupported compression formats are read as plain text.

Attempting to read such files does not generate a diagnostic.

Resolution: Fixed in 0.2

Impala server raises a null pointer exception when running an HBase query.

When querying an HBase table whose row-key is string type, the Impala server may raise a null pointer exception.

Resolution: Fixed in 0.2

Apache Kudu Fixed Issues

This topic includes the issues fixed in all generally available versions (GA) of Apache Kudu.

Issues Fixed in Kudu for CDH 5.16.2

The following is a list of upstream issues fixed in Kudu for CDH 5.16.2. For the full list of fixed issues for all CDH components in CDH 5.16.2, see [Upstream Issues Fixed in CDH 5.16.2](#).

- [KUDU-1678](#) - Fixed a crash caused by a race condition between altering tablet schemas and deleting tablet replicas.
- [KUDU-2195](#) - Now you can use the `--cmeta_force_fsync` flag to fsync Kudu's consensus metadata more aggressively. Setting this to `true` may decrease Kudu's performance, but will improve its durability in the face of power failures and forced shutdowns. The issue was much more likely to happen when Kudu was running on XFS.
- [KUDU-2463](#) - Fixed an issue in which incorrect results would be occasionally returned in scans following a server restart.
- [CDH-76920](#) - Fixed the issue where the Kudu CLI crashes when running the 'kudu cluster rebalance' sub-command on some platforms.

Issues Fixed in Kudu 1.7.0 / CDH 5.16.1

The following is a list of upstream issues fixed in Kudu 1.7.0 / CDH 5.16.1. For the full list of fixed issues for all CDH components in CDH 5.16.1, see [Upstream Issues Fixed in CDH 5.16.1](#).

- [KUDU-2260](#) - Fixed a rare issue where system failure could leave unexpected null bytes at the end of metadata files, causing Kudu to be unable to restart.
- [KUDU-2364](#) - Fixed an issue when a tablet server was wiped and recreated with the same RPC address, ksck listed it twice, both as healthy, even though only one of them was there.
- [KUDU-2412](#) - The kudu-python client can now compile in environments where `__int128` is not supported. This was most commonly el6 environments.
- [KUDU-2509](#) - Fixed an issue that might result in a crash of a tablet server in case of a WAL replay error while bootstrapping a tablet.
- [KUDU-2580](#) - Fixed authentication token reacquisition in the C++ client.
- Fixed an issue that caused the kudu CLI tool to unexpectedly exit when the connection to the master or tserver was abruptly closed.

Issues Fixed in Kudu 1.7.0 / CDH 5.15.2

The following is a list of upstream issues fixed in Kudu 1.7.0 / CDH 5.15.2. For the full list of fixed issues for all CDH components in CDH 5.15.2, see [Upstream Issues Fixed in CDH 5.16.1](#).

- [KUDU-2463](#) - Fixed an issue in which incorrect results would be returned in scans following a server restart.
- [KUDU-2509](#) - Fixed an issue that might result in a crash of a tablet server in case of a WAL replay error while bootstrapping a tablet.
- [KUDU-2580](#) - Fixed authentication token reacquisition in the C++ client.

Issues Fixed in Kudu 1.7.0 / CDH 5.15.1

The following is a list of upstream issues fixed in Kudu 1.7.0 / CDH 5.15.1. For the full list of fixed issues for all CDH components in CDH 5.15.1, see [Upstream Issues Fixed in CDH 5.15.1](#).

- [KUDU-2367](#) - Fixed an issue where a permanently failed tablet replica was not properly identified, which could cause the tablet not to re-replicate in very small clusters.
- [KUDU-2377](#) - Fixed an issue that caused Kudu servers to fail to start when `RLIMIT_NPROC=-1`.
- [KUDU-2378](#) - Fixed unaligned loads of `int128` from rows.

- [KUDU-2379](#) - Fixed an issue that caused secure Spark jobs to fail.
- [KUDU-2416](#) - Fixed PartialRow.setMin.
- [KUDU-2443](#) - Fixed replica movement and replacement for RF=1.
- [KUDU-2447](#) - Fixed the tablet server crash with the error, "NONE predicate can not be pushed into key".
- [KUDU-2478](#) - Restored Python 2.6 compatibility.
- Added the ability to adjust scan timeouts in Spark.
- Increased the timeout to begin tablet copies, which improves Kudu's re-replication time when the cluster is busy.
- Fixed a NullPointerException thrown when calling ColumnSchema#toString on non-decimal types.
- Greatly improved the performance of many types of queries on tables from which many rows have been deleted.
- Fixed an issue that caused partition pruning to be too conservative for queries from the Java client that use Decimal predicates.

Issues Fixed in Kudu 1.7.0 / CDH 5.15.0

The following is a list of upstream issues fixed in Kudu 1.7.0 / CDH 5.15.0. For the full list of fixed issues for all CDH components in CDH 5.15.0, see [Upstream Issues Fixed in CDH 5.15.0](#).

- [KUDU-1613](#) - Fixed a scenario where the on-disk data of a tablet server was completely erased and a new tablet server was started on the same host. This issue could prevent tablet replicas previously hosted on the server from being evicted and re-replicated. Tablets now immediately evict replicas that respond with a different server UUID than expected.
- [KUDU-1927](#) - Fixed a rare race condition when connecting to masters during their startup which might cause a client to get a response without a CA certificate and/or authentication token. This would cause the client to fail to authenticate with other servers in the cluster. The leader master now always sends a CA certificate and an authentication token (when applicable) to a Kudu client with a successful ConnectToMaster response.
- [KUDU-2262](#) - The Kudu Java client now will retry a connection if no master is discovered as a leader, and the user has a valid authentication token. This avoids failure in recoverable cases when masters are in the process of the very first leader election after starting up.
- [KUDU-2264](#) - The Java client will now automatically attempt to re-acquire Kerberos credentials from the ticket cache when the prior credentials are about to expire. This allows client instances to persist longer than the expiration time of a single Kerberos ticket so long as some other process renews the credentials in the ticket cache. Documentation on interacting with Kerberos authentication has been added to the Javadoc for the `AsyncKuduClient` class.
- [KUDU-2265](#) - Follower masters are now able to verify authentication tokens even if they have never been a leader. Prior to this fix, if a follower master had never been a leader, clients would be unable to authenticate to that master, resulting in spurious error messages being logged.
- [KUDU-2295](#) - Fixed a tablet server crash when a tablet replica is deleted during a scan.
- [KUDU-2312](#) - The evaluation order of predicates in scans with multiple predicates has been made deterministic. Due to a bug, this was not necessarily the case previously. Predicates are applied in most to least selective order, with ties broken by column index. The evaluation order may change in the future, particularly when better column statistics are made available internally.
- [KUDU-2331](#) - Previously, the `kudu tablet change_config move_replica` tool required all tablet servers in the cluster to be available when performing a move. This restriction has been relaxed: only the tablet server that will receive a replica of the tablet being moved and the hosts of the tablet's existing replicas need to be available for the move to occur.
- [KUDU-2343](#) - Fixed a bug in the Java client which prevented the client from locating the new leader master after a leader failover in the case that the previous leader either remained online or restarted quickly. This bug resulted in the client timing out operations with errors indicating that there was no leader master.
- [KUDU-2259](#) - The Unix process username of the client is now included inside the exported security credentials, so that the effective username of clients who import credentials and subsequently use unauthenticated (SASL PLAIN) connections matches the client who exported the security credentials. For example, this is useful to let the Spark executors know which username to use if the Spark driver has no authentication token. This change only affects clusters with encryption disabled using `--rpc-encryption=disabled`.

Issues Fixed in Kudu for CDH 5.14.4

The following is a list of upstream issues fixed in Kudu for CDH 5.14.4. For the full list of fixed issues for all CDH components in CDH 5.14.4, see [Upstream Issues Fixed in CDH 5.14.4](#).

- [KUDU-2331](#) - The `kudu tablet change_config move_replica` tool no longer fails for an unavailable tablet server if the tablet server is not the source or destination of the move.
- Improved the speed of scans in cases where a large number of rows have previously been deleted from the table.

Issues Fixed in Kudu 1.6.0 / CDH 5.14.2

The following is a list of upstream issues fixed in Kudu 1.6.0 / CDH 5.14.2. For the full list of fixed issues for all CDH components in CDH 5.14.2, see [Upstream Issues Fixed in CDH 5.14.2](#).

- [KUDU-1613](#) - Fixed an issue where a reformatted server could prevent re-replication of the replicas it previously hosted.
- [KUDU-2238](#) - Fixed an issue where large updates were not flushed to disk under memory pressure.
- [KUDU-2251](#) - Fixed a crash due to an overflow bug that could occur when a tablet was updated very frequently.
- [KUDU-2274](#) - Fixed an issue where copying over a tombstoned replica could produce a replica in an inconsistent state.
- [KUDU-2343](#) - Fixed an issue where the Java client would fail to connect to the leader master after a leadership change.

Issues Fixed in Kudu 1.6.0 / CDH 5.14.0

The following is a list of upstream issues fixed in Kudu 1.6.0 / CDH 5.14.0. For the full list of fixed issues for all CDH components in CDH 5.14.0, see [Upstream Issues Fixed in CDH 5.14.0](#).

- [KUDU-1078](#) - Fixed an error message commonly found in tablet server logs indicating that operations were being read "from the future".
- [KUDU-1411](#) - HybridTime timestamp propagation now works in the Java client when using scan tokens.
- [KUDU-2044](#) - Tombstoned tablets no longer report metrics.
- [KUDU-2173](#) - Fixed a bug in the C++ client which could cause tablets to be erroneously pruned, or skipped, during certain scans, resulting in fewer results than expected being returned from queries. The bug only affected tables whose range partition columns are a proper prefix of the primary key.
- [KUDU-2188](#) - Published Kudu Java artifacts are now fully compatible with JRE 7 and JRE 8. There was previously a bug in the release process which made them compatible only with JRE 8.
- Fixed a typo in the list of default TLS ciphers used by Kudu servers. As a result, two additional cipher suites are now available:
 - ECDHE-RSA-AES128-SHA256 TLSv1.2 Kx=ECDH Au=RSA Enc=AES(128) Mac=SHA256
 - AES256-GCM-SHA384 TLSv1.2 Kx=RSA Au=RSA Enc=AESGCM(256) Mac=AEAD
- [KUDU-2231](#) - Sparse column predicates no longer cause excessive data-block reads.

Issues Fixed in Kudu 1.6.0 / CDH 5.13.3

This is a complete list of upstream issues fixed in Kudu 1.6.0 / CDH 5.13.3. For the full list of fixed issues for all CDH components in CDH 5.13.3, see [Upstream Issues Fixed in CDH 5.13.3](#).

- [KUDU-2238](#) - Fixed a bug where large updates were not flushed to disk, even when the server is under memory pressure..
- [KUDU-2274](#) - Fixed a very rare bug where a tombstone being copied over could produce a tablet replica in an inconsistent state.
- [KUDU-2343](#) - Java client doesn't properly reconnect to leader master when old leader is online.

Issues Fixed in Kudu 1.6.0 / CDH 5.13.2

This is a complete list of upstream issues fixed in Kudu 1.6.0 / CDH 5.13.2. For the full list of fixed issues for all CDH components in CDH 5.13.2, see [Upstream Issues Fixed in CDH 5.13.2](#).

- [KUDU-2147](#) - Consensus: Unknown leader should not be treated as valid UUID
- [KUDU-2231](#) - Sparse column predicate can cause excessive data-block reads.
- [KUDU-2251](#) - rowset size can overflow int in RowSetInfo.
- rpc: Allow setting `--rpc_tls_min_protocol` on older RHEL versions.

Issues Fixed in Kudu 1.5.0 / CDH 5.13.1

This is a complete list of upstream issues fixed in Kudu 1.5.0 / CDH 5.13.1. For the full list of fixed issues for all CDH components in CDH 5.13.1, see [Upstream Issues Fixed in CDH 5.13.1](#).

- [KUDU-1788](#) - Increase Raft RPC timeout to 30sec to avoid fruitless retries.
- [KUDU-2130](#) - (part 2): more fixes for ITClientStress
- [KUDU-2130](#) - java client: handle termination during negotiation edge case
- [KUDU-2167](#) - fix C++ client crash due to bad assumption regarding scan data
- [KUDU-2170](#) - Masters can start with duplicates specified
- [KUDU-2173](#) - Partitions are incorrectly pruned when range-partitioned on a PK prefix
- [KUDU-2188](#) - restore Java 7 compatibility to artifacts built with JDK8
- [KUDU-2209](#) - HybridClock doesn't handle changes in STA_NANO flag

Issues Fixed in Kudu 1.5.0 / CDH 5.13.0

- The Java Kudu client now automatically requests new authentication tokens after expiration. As a result, long-lived Java clients are now supported. See [KUDU-2013](#) for more details.
- Multiple Kerberos compatibility bugs have been fixed, including support for environments with disabled reverse DNS, FreeIPA compatibility, principal names including uppercase characters, and hosts without a FQDN.
- A bug in the binary prefix decoder which could cause a tablet server 'check' assertion crash has been fixed. The crash could only be triggered in very specific scenarios; see [KUDU-2085](#) for additional details.
- This is a complete list of upstream issues fixed in Kudu 1.5.0 / CDH 5.13.0. For the full list of fixed issues for all CDH components in CDH 5.13, see [Upstream Issues Fixed in CDH 5.13](#).
 - [KUDU-871](#) - Allow tombstoned tablets to vote
 - [KUDU-1125](#) - (part 1) catalog_manager: try to avoid unnecessarily rewriting tablet info
 - [KUDU-1407](#) - reassign failed tablets
 - [KUDU-1442](#) - log number of open log block containers
 - [KUDU-1544](#) - Race in Java client's AsyncKuduSession.apply()
 - [KUDU-1726](#) - Avoid fsync-per-block in tablet copy
 - [KUDU-1755](#) - Part 1: Improve tablet on disk size metric
 - [KUDU-1811](#) - C++ client: use larger batches when fetching scan tokens
 - [KUDU-1863](#) - improve overall safety of graceful server shutdown
 - [KUDU-1865](#) - Avoid heap allocation for payload slices
 - [KUDU-1894](#) - fixed deadlock in client.Connection
 - [KUDU-1911](#) - improve missing required arg message
 - [KUDU-1929](#) - [rpc] Allow using encrypted private keys for TLS
 - [KUDU-1942](#) - Kerberos fails to log in on hostnames with capital letters
 - [KUDU-1943](#) - Add BlockTransaction to Block Manager
 - [KUDU-1952](#) - Remove round-robin for block placement
 - [KUDU-1955](#) - refuse to use world-readable keytabs
 - [KUDU-2004](#) - Undefined behavior in TlsSocket::Writev()
 - [KUDU-2013](#) - Support for long lived auth tokens in Java client
 - [KUDU-2032](#) - Kerberos authentication fails with rdns disabled in krb5.conf
 - [KUDU-2039](#) - Fix the table count in the /tables page of master webUI
 - [KUDU-2041](#) - Fix negotiation deadlock
 - [KUDU-2049](#) - Fix too-strict CHECK in RleIntBlockDecoder::SeekToPositionInBlock

- [KUDU-2053](#) - Fix race in Java RequestTracker
- [KUDU-2058](#) - Fix LocatedTablet string comparisons
- [KUDU-2060](#) - Show primary keys in the master's table web UI page
- [KUDU-2066](#) - Add experimental Gradle build support
- [KUDU-2067](#) - Enable cfile checksumming by default
- [KUDU-2072](#) - upgrade to cmake 3.9.0 breaks sles12sp0 cmake patch
- [KUDU-2078](#) - Sink failure if batch size > session's flush buffer size
- [KUDU-2083](#) - Decrement running maintenance ops on failed prepare
- [KUDU-2085](#) - Fix crash when seeking past end of prefix-encoded blocks
- [KUDU-2087](#) - Fix failure to map Kerberos principal to username with FreeIPA
- [KUDU-2088](#) - Synchronizer may not go out of scope with outstanding references
- [KUDU-2091](#) - Certificates with intermediate CA's do not work with Kudu
- [KUDU-2101](#) - Include a table summary at the bottom
- [KUDU-2102](#) - fix PosixRWFile::Sync to guarantee durability when used concurrently
- [KUDU-2103](#) - [java]Canonicalize hostnames in client
- [KUDU-2104](#) - Upgrade to Spark 2.2.0
- [KUDU-2114](#) - Don't re-delete tombstoned replicas
- [KUDU-2118](#) - Fully shut down TabletReplica on delete
- [KUDU-2123](#) - Auto-vivify cmeta on tombstoned replicas if doesn't exist at startup
- [KUDU-2131](#) - switch to LIFO log container retrieval
- [KUDU-2138](#) - delete failed replicas in tablet report
- [KUDU-2141](#) - master: Remove DCHECK when tablet report has no opid_index

Issues Fixed in Kudu 1.4.0 / CDH 5.12.2

Apache Kudu 1.4.0 / CDH 5.12.2 is a bug-fix release which includes the following fixes:

- [KUDU-2209](#) - Fixed Kudu's handling of the STA_NANO flag, which caused spurious crashes due to one Kudu node thinking another node had a timestamp from the future.
- [KUDU-1788](#) - Increased the Raft RPC timeout to 30 seconds to avoid unsuccessful retries.
- [KUDU-2170](#) - Disallowed specifying the same address multiple times in Kudu's master list.
- [KUDU-2083](#) - Fixed an issue where, in rare circumstances, Kudu reduced the amount of maintenance ops it could run concurrently. The issue could cause the Kudu process to run out of memory and crash.
- [KUDU-1942](#) - Now Kerberos can successfully log in on hostnames with capital letters.
- [KUDU-2173](#) - Fixed an issue where scans with a predicate on a prefix of a range partition key could fail to return matching values.
- [KUDU-2167](#) - Fixed an issue where the C++ client would crash when the server filtered out all data in a scan RPC.
- [KUDU-2032](#) - Kerberos authentication no longer fails when `rdns = false` is configured in `krb5.conf`.
- Fixed a segmentation fault that could occur when running the `ksck` command against a master soon after the master was started.

Issues Fixed in Kudu 1.4.0 / CDH 5.12.1

Apache Kudu 1.4.0 / CDH 5.12.1 is a bug-fix release which includes the following fixes:

- [KUDU-2085](#) - Fixed a bug that caused crashes when seeking past the end of prefix-encoded blocks.
- [KUDU-2087](#) - Fixed an issue where Kudu would fail to start when Kerberos was enabled in FreeIPA-configured deployments.
- [KUDU-2053](#) - Fixed a race condition in the Java RequestTracker.
- [KUDU-2049](#) - Fixed an issue where scans on RLE-encoded integer columns would sometimes cause CHECK failures due to the CHECK condition being too strict.
- [KUDU-2052](#) - Kudu now uses `XFS_IOC_UNRESVSP64 ioctl` to punch holes on xfs filesystems. This fixes an issue with slow startup times on xfs when hole punching was done via `fallocate()`.
- [KUDU-1956](#) - Kudu will no longer crash if faced with a race condition when selecting rowsets for compaction.

Issues Fixed in Kudu 1.4.0 / CDH 5.12.0

- [KUDU-2020](#) - Fixed an issue where re-replication after a failure would proceed significantly slower than expected. This bug caused many tablets to be unnecessarily copied multiple times before successfully being considered re-replicated, resulting in significantly more network and IO bandwidth usage than expected. Mean time to recovery on clusters with large amounts of data is improved by up to 10x by this fix.
- [KUDU-1982](#) - Fixed an issue where the Java client would call `NetworkInterface.getByInetAddress` very often, causing performance problems particularly on Windows where this function can be quite slow.
- [KUDU-1755](#) - Improved the accuracy of the `on_disk_size` replica metrics to include the size consumed by bloom filters, primary key indexes, superblock metadata, and delta files. Note that because the size metric is now more accurate, the reported values are expected to increase after upgrading to Kudu 1.4.0. This does not indicate that replicas are using more space after the upgrade; rather, it is now accurately reporting the amount of space that has always been used.
- [KUDU-1192](#) - Kudu servers will now periodically flush their log messages to disk even if no `WARNING`-level messages have been logged. This makes it easier to tail the logs to see progress output during normal startup.

Issues Fixed in Kudu 1.3.0 / CDH 5.11.2

Apache Kudu 1.3.0 / CDH 5.11.2 is a bug-fix release which includes the following fixes:

- [KUDU-2053](#) - Fixed a race condition in the Java `RequestTracker`.
- [KUDU-2049](#) - Fixed an issue where scans on RLE-encoded integer columns would sometimes cause CHECK failures due to the CHECK condition being too strict.
- [KUDU-1963](#) - Fixed an issue where the Java client misdiagnoses an error and logs a `NullPointerException` when a connection is closed by client while a negotiation is in progress.
- [KUDU-1853](#) - Fixed an issue where data blocks could be orphaned after a failed tablet copy.

Issues Fixed in Kudu 1.3.0 / CDH 5.11.1

Apache Kudu 1.3.0 / CDH 5.11.1 is a bug-fix release which includes the following fixes:

- [KUDU-1999](#) - Fixed an issue where the Kudu Spark connector would fail to kinit with the principal and keytab provided to a job.
- [KUDU-1993](#) - Fixed a validation issue with grouped gflags.
- [KUDU-1981](#) - Fixed an issue where Kudu server components would fail to start on machines with fully-qualified domain names longer than 64 characters when security was enabled. This was due to hard-coded restrictions in the OpenSSL library.
- [KUDU-1607](#) - Fixed a case in which a tablet replica on a tablet server could retain blocks of data which prevented it from being fully deleted.
- [KUDU-1933](#) - Fixed an issue in which a tablet server would crash and fail to restart after a single tablet received more than two billion write operations.
- [KUDU-1964](#) - Fixed a performance degradation issue caused by OpenSSL locks under high concurrency.

Issues Fixed in Kudu 1.3.0 / CDH 5.11.0

- [KUDU-1968](#) - Fixed an issue in which the tablet server would delete an incorrect set of data blocks after an aborted attempt to copy a tablet from another server. This would produce data loss in unrelated tablets.
- [KUDU-1962](#) - Fixed a `NullPointerException` in the Java client in the case that the Kudu master is overloaded at the time the client requests location information. This could cause client applications to hang indefinitely regardless of configured timeouts.
- [KUDU-1893](#) - Fixed a critical bug in which wrong results would be returned when evaluating predicates applied to columns added using the `ALTER TABLE` operation.
- [KUDU-1905](#) - Fixed an issue where Kudu would crash after reinserts that resulted in an empty change list. This occurred in cases where the primary key was composed of all columns.

- [KUDU-1899](#) - Fixed a crash that occurred after inserting a row with an empty string as the single-column primary key.
- [KUDU-1904](#) - Fixed a potential crash when performing random reads against a column using RLE encoding and containing long runs of NULL values.
- [KUDU-1856](#) - Fixed an issue in which disk space could be leaked by Kudu servers storing data on partitions using the XFS file system. Any leaked disk space will now be automatically recovered upon upgrade.
- [KUDU-1888](#), [KUDU-1906](#) - Fixed multiple issues in the Java client where operation callbacks would never be triggered, causing the client to hang.

Issues Fixed in Kudu 1.2.0 / CDH 5.10.2

Apache Kudu 1.2.x / CDH 5.10.2 includes the following fixed issues.

- [KUDU-1933](#) - Fixed an issue that truncated the 64-bit log index in the OpId to 32 bits, causing overflow of the log index.
- [KUDU-1607](#) - Fixed an issue where Kudu could not delete failed tablets using the `DROP TABLE` command.
- [KUDU-1905](#) - Allow reinserts on tables when all columns are part of the primary key.
- [KUDU-1893](#) - Made a fix to avoid incorrect NULL results and ensure evaluation of predicates for columns added after table creation.

Issues Fixed in Kudu 1.2.0 / CDH 5.10.1

Apache Kudu 1.2.x / CDH 5.10.1 includes the following fixed issues.

- [KUDU-1904](#) - Fixed a bug where RLE columns with only NULL values would crash on scan.
- [KUDU-1899](#) - Fixed an issue where tablet servers would crash after inserting an empty string primary key (" ").
- [KUDU-1851](#) - Fixed an issue with the Python client which would crash whenever a `TableAlterer` is instantiated directly.
- [KUDU-1852](#) - `KuduTableAlterer` will no longer crash when given `nullptr` range bound arguments.
- [KUDU-1821](#) - Improved warnings when the catalog manager starts.

Issues Fixed in Kudu 1.2.0 / CDH 5.10.0

See [Issues resolved for Kudu 1.2.0](#) and [Git changes between 1.1.x and 1.2.x](#).

- [KUDU-1508](#) - Fixed a long-standing issue in which running Kudu on `ext4` file systems could cause file system corruption. While this issue has been known to still manifest in certain rare cases, the corruption is harmless and can be repaired as part of a regular `fsck`. Switching from `ext4` to `xfs` will also solve the problem.
- [KUDU-1399](#) - Implemented an LRU cache for open files, which prevents running out of file descriptors on long-lived Kudu clusters. By default, Kudu will limit its file descriptor usage to half of its configured `ulimit`.
- [Gerrit #5192](#) - Fixed an issue which caused data corruption and crashes in the case that a table had a non-composite (single-column) primary key, and that column was specified to use `DICT_ENCODING` or `BITSHUFFLE` encodings. If a table with an affected schema was written in previous versions of Kudu, the corruption will not be automatically repaired; users are encouraged to re-insert such tables after upgrading to Kudu 1.2 or later.
- [Gerrit #5541](#) - Fixed a bug in the Spark `KuduRDD` implementation which could cause rows in the result set to be silently skipped in some cases.
- [KUDU-1551](#) - Fixed an issue in which the tablet server would crash on restart in the case that it had previously crashed during the process of allocating a new WAL segment.
- [KUDU-1764](#) - Fixed an issue where Kudu servers would leak approximately 16-32MB of disk space for every 10GB of data written to disk. After upgrading to Kudu 1.2 or later, any disk space leaked in previous versions will be automatically recovered on startup.
- [KUDU-1750](#) - Fixed an issue where the API to drop a range partition would drop any partition with a matching lower `_or_` upper bound, rather than any partition with matching lower `_and_` upper bound.
- [KUDU-1766](#) - Fixed an issue in the Java client where equality predicates which compared an integer column to its maximum possible value (e.g. `Integer.MAX_VALUE`) would return incorrect results.

- [KUDU-1780](#) - Fixed the kudu-client Java artifact to properly shade classes in the `com.google.thirdparty` namespace. The lack of proper shading in prior releases could cause conflicts with certain versions of Google Guava.
- [Gerrit #5327](#) - Fixed shading issues in the `kudu-flume-sink` Java artifact. The sink now expects that Hadoop dependencies are provided by Flume, and properly shades the Kudu client's dependencies.
- Fixed a few issues using the Python client library from Python 3.

Cloudera Manager 5 Release Notes

These Release Notes provide information on the new features and known issues and limitations for Cloudera Manager 5. These Release Notes also include fixed issues for releases starting from Cloudera Manager 5.0.0 beta 1.

To view the Release Notes (or other documentation) for a specific Cloudera Manager release, go to [Cloudera Documentation](#), click a major version link, and use the drop-down menu to select the release.

For information about supported operating systems, and other requirements for using Cloudera Manager, see [CDH 5 and Cloudera Manager 5 Requirements and Supported Versions](#) on page 20.

New Features and Changes in Cloudera Manager 5

The following sections describe what's new and changed in each Cloudera Manager 5 release.

What's New in Cloudera Manager 5

The following sections describe what is new in each Cloudera Manager 5 release.

What's New in Cloudera Manager 5.16.2

Cloudera Manager 5.16.2 is a maintenance release with several fixed issues. See [Issues Fixed in Cloudera Manager 5.16.2](#) on page 2549.

What's New in Cloudera Manager 5.16.1

Cloudera Manager 5.16.1 is the first release of Cloudera Manager 5.16 and also has several fixed issues. See [Issues Fixed in Cloudera Manager 5.16.1](#) on page 2549.

OpenJDK Support

OpenJDK is now supported with Cloudera Manager and CDH 5.16. For information about migrating your Oracle JDK to OpenJDK, see [Upgrading the JDK](#). To use OpenJDK in a new deployment, see [Install Java Development Kit](#).

Cloudera Bug: CDH-60334

Backup and Disaster Recovery (BDR) Log retention

You can now configure the number of days Cloudera Manager retains BDR logs with the **Backup and Disaster Log Retention** property. For more information, see [BDR Log Retention](#).

Impala

Revised Impala charts library

The Impala chart library has been updated to include more meaningful metrics and remove rarely used plots.

Cloudera Bug: OPSAPS-46020

Impala Daemon's JVM heap size can now be configured

The Impala Daemon's JVM heap size can now be configured now using the **Java Heap Size of Impala Daemon in Bytes** configuration property. The property defaults to 4GB, and, like all memory parameters, may require tuning.

Cloudera Bug: OPSAPS-41238

New Impala Alert when reaching maximum client connections

A new health check has been added to report and alert when the Impala Daemon is reaching the maximum capacity of concurrent client connections. The **Impala Daemon Max Client Connections** configuration parameter has been added to configure this value.

Cloudera Bug: OPSAPS-46025

New Impala Idle Query Timeout and Idle Session Timeout configuration properties

The following new configuration properties have been added for Impala: **Idle Query Timeout** and **Idle Session Timeout** parameters.

Cloudera Bug: OPSAPS-38917

Impala Assignment Locality health test removed from Cloudera Manager

The **Impala Assignment Locality** health test has been removed.

Cloudera Bug: OPSAPS-46807

The Solr collection statistics page was not shown for Auditor or Dashboard user roles

A Cloudera Manager Dashboard or Audit user can now see the **Solr Collection Statistics** and the **HBase Table Statistics** page.

Cloudera Bug: OPSAPS-45238, OPSAPS-45238

New Kafka Health Tests

Two new Kafka Broker Health Tests have been added to Cloudera Manager: **Kafka Broker Swap Memory Usage** and **Kafka Broker Unexpected Exits**. These health tests are enabled by default when Kafka is running as a service in Cloudera Manager version 5.14 and later. See [Kafka Broker Health Tests](#).

Cloudera Bug: OPSAPS-45002

Sentry Object Ownership configurations in Cloudera Manager

A new configuration parameter has been added to the Sentry Configuration page that enables Sentry OWNER privileges (this parameter is disabled by default). For more information, see [Setting Object Owner Privileges in Cloudera Manager](#).

Cloudera Bug: OPSAPS-47434

Downgrade to Cloudera Express disabled when using HDFS encryption at rest

Added a warning to the **Downgrade License** dialog box if the installation uses HDFS encryption with KMS Key Trustee server and disabled the downgrade in this case. Users are instructed to convert their data to an unencrypted format before proceeding.

Cloudera Bug: OPSAPS-44373

ZooKeeper configuration changes

The **Enable Kerberos Authentication** and **Enable Server to Server SASL Authentication** settings in ZooKeeper are now linked together. If either parameter is changed to on or off, the other parameter will automatically change to the same value. A warning message has been added to indicate the effect.

Cloudera Bug: OPSAPS-46628

What's New in Cloudera Manager 5.15.2

Cloudera Manager 5.15.2 is a maintenance release with several fixed issues. See [Issues Fixed in Cloudera Manager 5.15.2](#) on page 2552.

What's New in Cloudera Manager 5.15.1

Cloudera Manager 5.15.1 is a maintenance release with several fixed issues. See [Issues Fixed in Cloudera Manager 5.15.1](#) on page 2553.

Dynamic Resource Pools UI

The Dynamic Resource Pools user interface now supports the following placement rules and pool creation policy can be separately configured for the parent group as well as the individual user group:

- root.primaryGroup.username
- root.secondaryExistingGroup.username
- root.[pool name].username.

Previously, only the `create="true|false"` flag could be added to the inner element of the `nestedUserQueue` element. This meant that a `root.primaryGroup` or `root.secondaryExistingGroup` pool could be created, which was not correct. Now, you can add the `create="true|false"` flag to the actual `nestedUserQueue` element as well as the inner element of the `nestedUserQueue` element. An additional restriction is that if `root.<parent>.username` should use an existing pool (`create = false`), then `root.<parent>` must also use an existing pool.

New fields in Cloudera Manager Python API for BDR

Cloudera Manager now exposes the following fields in the Python API to match the fields returned by the REST API:

- `raiseSnapshotDiffFailures` to `ApiHdfsReplicationArguments`
- `remainingTime`, `throughput`, and `estimatedCompletionTime` to `ApiHdfsCloudReplicationArguments`
- `numThreads`, `runInvalidateMetadata` to `ApiHiveReplicationArguments`
- `displayName`, `description` to `ApiReplicationSchedule`

These fields were already exposed in the Java Cloudera Manager API.

What's New in Cloudera Manager 5.15.0

- **Backup and Disaster Recovery (BDR)**

- **ADLS** - You can now replicate HDFS files and Hive data to and from Microsoft ADLS. To use ADLS as the source or destination, you must add Azure credentials to Cloudera Manager.

Note that the lowest supported version for Cloudera Manager and CDH for BDR replication to/from ADLS is version 5.15.0.

- **Metrics** - The amount of data read and written through Amazon S3 and Microsoft ADLS storage by MapReduce jobs can now be viewed as cluster metrics, such as `s3a_bytes_read` and `adl_bytes_written`.
- **Multi-threaded import and export for Hive Replication** - You can now configure the number of threads used for import and export during Hive Replication. When you create or modify a Hive Replication schedule, configure the **Number of concurrent HMS connections** on the **Advanced** tab. Increasing the number of threads can improve BDR performance. By default, any new replication schedules will use 5 connections.

If you set the value to 1 or more, BDR uses multi-threading with the number of connections specified.

If you set the value to 0 or fewer, BDR uses single threading and a single connection.

The source and destination clusters must both run Cloudera Manager 5.15.0 or higher in order to use this feature.

- **Security** - To improve security, BDR now uses an encrypted Hadoop credential store to authenticate with cloud providers, such as Amazon S3 or Microsoft ADLS, when backing up or restoring HDFS or Hive data.
- **Statistics** - Hive Replication phases now show the number of Hive Objects found/processed. Each type of Hive object is represented separately: databases, tables, indexes, functions, partitions and column statistics. This information can be used to determine how many objects are replicated in each run. This can also be used to deduce how long will it take for Hive Replication to complete.
- **Snapshot diff-based replication** - This feature compares two HDFS snapshots to reduce the number of files scanned during the copy-listing phase of replication. It can speed up replication performance when large number of files are unchanged between replications.

You must enable immutable snapshots for HDFS to use snapshot diff-based replication for BDR.

This feature is on by default. You can configure replications to abort on snapshot diff failure when you create or edit a replication schedule.

See the following pages for guidelines on using snapshot diff-based replication: [guidelines for Hive replication](#) and [guidelines for HDFS replication](#).

- **HDFS** - You can now enable immutable snapshots for HDFS with Cloudera Manager. Enabling this feature enables snapshot diff-based copy listing for BDR.

In the Cloudera Manager Admin Console, navigate to **Clusters > <HDFS cluster> > Configuration** and search for **Enable Immutable Snapshots**.

- **Maintenance and Support**

- **Cluster Restart** - Improved cluster restart performance.
- **Kudu** - Cloudera Manager now supports gathering the output of the ksck diagnostic tool from Kudu. This output is now collected in the diagnostic bundles.

- **Impala**

- **File Handle Cache** - You can now configure and monitor the following Impala parameters using Cloudera Manager: `max_cached_file_handles` and `unused_file_handle_timeout_sec`.
- **KRPC Port** - You can now configure the `krpc_port` startup parameter using Cloudera Manager. The default value is 27000
- **Metrics** - Cloudera Manager now collects the following metrics for Impala: `impala_jvm_heap_committed_usage_bytes`, `impala_jvm_heap_current_usage_bytes`, `impala_jvm_heap_init_usage_bytes`, `impala_jvm_heap_max_usage_bytes`. Impala administrators may find these metrics useful for monitoring Catalog size and Impala Daemon health and the amount of memory used by the JVM embedded in the Impala Daemon process. These are useful for understanding memory consumption, particularly the memory consumption of the Catalog cache stored in coordinator Impala Daemons.

- **Parcels** - When configuration changes are made on the **Parcels > Configuration** page, Cloudera Manager now automatically checks for any new parcels. Additionally, you can disable recurring checks for parcels by setting the **Parcel Update Frequency** to '0' on the **Parcels > Configurations** page.

- **Upgrade**

- **Agents** -

Cloudera Manager Agent upgrade now works with mixed environments where there are different distribution/versions of operating systems in a single Cloudera Manager deployment. These agents can be upgraded as a group for each set of operating systems.

The agents are grouped and displayed on a new page in the Cloudera Manager Upgrade wizard.

- **Documentation** - You can now find a link to the latest upgrade documentation in the Cloudera Manager Admin Console by going to **Support > Upgrade Guide**.

Upgrade documentation now includes new interactive features that allow you to select your operating system, upgrade version, database type, CDH installation type (Parcels or Packages), and other features, and a customized page displays only the steps required for your upgrade. See:

- [Upgrading Cloudera Manager Using Packages](#)
- [Upgrading CDH](#)

- **Summary Page** - All potential issues, conflicts, action items, and pre-upgrade checks are summarized on the first page of the CDH Upgrade wizard.

What's New in Cloudera Manager 5.14.4

Cloudera Manager 5.14.4 is a maintenance release with several fixed issues. See [Issues Fixed in Cloudera Manager 5.14.4](#) on page 2559.

What's New in Cloudera Manager 5.14.3

Cloudera Manager 5.14.3 is a maintenance release with several fixed issues. See [Issues Fixed in Cloudera Manager 5.14.3](#) on page 2561.

What's New in Cloudera Manager 5.14.2

Cloudera Manager 5.14.2 is a maintenance release with several fixed issues. See [Issues Fixed in Cloudera Manager 5.14.2](#) on page 2561.

What's New in Cloudera Manager 5.14.1

Cloudera Manager 5.14.1 is a maintenance release with several fixed issues. There is no corresponding CDH 5.14.1 release. See [Issues Fixed in Cloudera Manager 5.14.1](#) on page 2562.

What's New in Cloudera Manager 5.14.0

- **ADLS**

- You can now use Cloudera Manager to configure Microsoft Azure credentials for cluster access to ADLS. This access is enabled for running Hive and Impala queries on tables backed by data stored in ADLS and to browse ADLS data using Hue. See [Configuring ADLS Access Using Cloudera Manager](#).
- ADLS integration with Hue added. Hue can be setup to enable browsing files in Azure Data Lake Store, and Hue users can directly query and store data in ADLS without any intermediate moving or copying of data to and from HDFS.

- **Backup and Disaster Recovery (BDR)**

- Added the option to **Skip Checksum on Listing**. This option skips the checksum comparison between two files to determine whether they are the same or not. To detect modifications, BDR will use the file size and last modified time instead. Skipping the checksum on listing can improve performance.
- You can now perform a rolling restart of clusters or services such as HDFS or Hive while replication jobs are running.
- On the **Replication Schedules** page, you can see the **Throughput** and **Progress** for replications that are running. Additionally, this information can be seen in the log files.
- Added Advanced Configuration Snippets for `hdfs-site.xml` and `core-site.xml` that can be used for HDFS replication. You can use these to tune behavior such as the HDFS replication factor of files written by BDR. Previously, you had to use the HDFS Client Advanced Configuration Snippet, which would affect all HDFS client configuration on the target cluster.

- **Custom Service Descriptor (CSD)**

- Added a new capability for custom services. CSD authors can now specify, in the topology descriptor for a role type, that only an odd number of instances should be running.
- The CSD versioning logic has been improved. When a CSD loads during Cloudera Manager startup, only the latest version of a CSD service type is loaded. The version number has precedence over the CDH compatibility range.

- **Diagnostic Bundles**

Diagnostic Bundles now include the following additional data: cluster utilization report data and Cloudera Navigator dashboard data (if the cluster has a Cloudera Navigator instance).

- **Health Tests and Monitoring**

- Added a new health test for Linux host entropy monitoring. By default, the health test issues a warning for entropy < 100 and an alert for entropy < 50. Additionally, a chart is available in the Charts Library.
- Improved Kafka health monitoring:
 - Cloudera Manager now aggregates Kafka Broker health to the health of the Kafka service. If greater than 5% percent of the brokers' health statuses are concerning, the Kafka service's health will be set to concerning. If 50% or more of the brokers' health statuses are bad, then the service's health will be set to bad.
 - If at least one partition hosted by a broker is offline, the service's health will now be bad.

- If replication of any partitions hosted by a broker is lagging, then the health of the broker will be set to concerning.

- **Host Maintenance Mode**

Cloudera Manager now provides more options for performing host maintenance and decommissioning. In addition to decommissioning, you can now specify in the same dialog whether to suppress alerts from the decommissioned host and, for hosts running a DataNode role, you can specify whether or not to replicate under-replicated data blocks to other DataNodes during maintenance. This feature takes the DataNode into "offline" state and is useful when performing minor maintenance on hosts, such as adding memory or changing network cards or cables where the maintenance window is expected to be short, the disk drives on the host remain in place, and the extra cluster resources would be consumed in replicating blocks is undesirable. See [Tuning and Troubleshooting Host Decommissioning](#).

- **Installation**

- After a host installation is completed with the Cloudera Manager Admin Console or API, a command on the **Commands** page provides a link to a zipped file that contains the installation logs for each host.
- You can now add a service to a cluster that does not have any services. Previously, you could not add services to a cluster without deleting and re-adding the cluster if the cluster was empty.

- **Logs**

The `cloudera-scm-agent.out` directory now uses the same location as Agent log files. Specify the location in the Agent `config.ini` file.

- **Set default Oozie Configuration**

Added a new Oozie Advanced Configuration Snippet for `action-conf/default.xml`. Use this Advanced Configuration Snippet to propagate default configuration properties to all Oozie actions. [OPSAPS-42546](#)

- **Resource Pools**

- You can now specify a percentage for the maximum CPU and memory resources for a YARN resource pool. Previously, you could only specify a static value.
- Dynamic resource pools can be sorted on the configuration page, status page, and utilization reports page in the Cloudera Manager Admin Console.

- **Search and Sentry**

The Keystore Indexer (HBase indexer) can be configured to use Sentry service for authorization.

What's New in Cloudera Manager 5.13.3

Cloudera Manager 5.13.3 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.13.3](#) on page 2568.

What's New in Cloudera Manager 5.13.2

Cloudera Manager 5.13.2 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.13.2](#) on page 2568.

What's New in Cloudera Manager 5.13.1

Cloudera Manager 5.13.1 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.13.1](#) on page 2570.

What's New in Cloudera Manager 5.13.0

- **Cloudera Data Science Workbench**

Cloudera Data Science Workbench is now available as an add-on service for Cloudera Manager. To this end, Cloudera Data Science Workbench is now distributed in a parcel that integrates with Cloudera Manager using a

Custom Service Descriptor (CSD). You can use Cloudera Manager to install, upgrade, and monitor Cloudera Data Science Workbench. Additionally, diagnostic data bundles for Cloudera Data Science Workbench can be generated and submitted to Cloudera through Cloudera Manager.

- **Dashboard User Role**

Users assigned the Dashboard User can perform the following actions:

- Create, edit, or remove their own dashboards
- Create new charts or add existing charts to their own dashboards
- View data in Cloudera Manager
- View service and monitoring information

- **Impala Query Profiles Time Display**

Impala query profiles downloaded from Cloudera Manager in text format now include a human-readable version for the value of each profile counter instead of the raw value. For example, for time counters before CM 5.13 only the raw nanoseconds value was shown: TotalTime: 492626971556. Now a human-readable value is shown alongside the raw value: TotalTime: 8.2m (492626971556).

- **Sentry High Availability**

The Sentry Service now supports High Availability. See [Sentry High Availability](#).

- **Licensing Management Improvements**

- **License Information**

A banner now indicates when there are 60, 30, 14, or 0 days left on a license.

- **Downgrading from Cloudera Enterprise to Cloudera Express**

Previously, downgrading from Enterprise to Express required editing the Cloudera Manager database. Now this task can be performed by the user by using the **Downgrade** button on the license page.

- **Enable Kerberos Wizard shows warning for hostnames with uppercase letters**

Because Kerberos principal names cannot include upper-case letters, the Enable Kerberos wizard welcome screen will show a warning if any hostnames containing uppercase characters are detected. However, the user will be able to continue with the wizard regardless of this warning. When the warning is shown, up to 10 such detected hostnames will be listed. If the list is longer than 10, the message also says only 10 are shown.

- **New Validations for Hadoop Configuration Properties User impersonation**

A new validation has been added for the various username configuration properties used by the Service Monitor to ensure that the properties are valid Linux usernames. The properties affected by this check are:

- HDFS User to Impersonate - HDFS service
- HBase User to Impersonate - HBase service
- MapReduce User to Impersonate - MapReduce service
- YARN Container Usage MapReduce Job User - YARN service

If these properties are not valid, the respective services will not start.

- **Additions to the Cloudera Manager API**

- **New Cloudera Manager API to control MaintenanceMode for Management Services**

Added REST API endpoints for Cloudera Management Service and Cloudera Management Service roles to enter and exit maintenance mode. See:

- **Enter Maintenance Mode for Cloudera Management**

Service:http://cloudera.github.io/cm_api/apidocs/v18/path_cm_service_commands_enterMaintenanceMode.html

- **Exit Maintenance Mode for Cloudera Management**

Service:http://cloudera.github.io/cm_api/apidocs/v18/path_cm_service_commands_exitMaintenanceMode.html

- **Enter Maintenance Mode for Cloudera Management Service**
roles:http://cloudera.github.io/cm_api/apidocs/v18/path_cm_service_roles_roleName_commands_enterMaintenanceMode.html
- **Exit Maintenance Mode for Cloudera Management Service**
roles:http://cloudera.github.io/cm_api/apidocs/v18/path_cm_service_roles_roleName_commands_exitMaintenanceMode.html

– New Cloudera Manager API for Cluster Utilization Report

Three new API endpoints have been added that return the [cluster utilization reports](#) in CSV format. See:

- **Cluster Utilization:**http://cloudera.github.io/cm_api/apidocs/v18/path_clusters_clusterName_utilization.html
- **Impala Utilization:**
http://cloudera.github.io/cm_api/apidocs/v18/path_clusters_clusterName_impalaUtilization.html
- **YARN Utilization:**
http://cloudera.github.io/cm_api/apidocs/v18/path_clusters_clusterName_yarnUtilization.html

See [Cloudera Manager API documentation](#).

– Delete credentials API at cluster level

There is a new API endpoint that allows you to delete all of the Kerberos credentials of services belonging to a single cluster, rather than all credentials generated by Cloudera Manager for all of its managed clusters. See:

http://cloudera.github.io/cm_api/apidocs/v18/path_clusters_clusterName_commands_deleteCredentials.html

• New Guardrails for Operating Cloudera Manager at Scale

– New Validator for Management Roles

A new validator warns when if multiple Management roles are running on the same host when Cloudera Manager manages more than 80 hosts.

– New Validators for Service Monitor and Host Monitor Memory Allocations

New validators have been added for Service and Host Monitor that give warnings if the heap and non-java memory for these roles are below the recommended values for the cluster. These validators depend on the size of the clusters, as well as on the types of services running in the cluster.

• Solr Chart Library

Chart Library for Solr now contains example charts for every metric available for the Solr Service.

• Improved BDR Performance

BDR replication performance has been improved by running the first phase of replication on the source cluster (the copy-listing phase, which creates a list of files and folders to be copied). This can dramatically improve performance in scenarios where there is high latency between the source and destination clusters. This feature requires Cloudera Manager 5.13 or higher on both the source and target cluster and can be disabled by setting a feature flag with the API.

• New Configuration Property for Descriptors

A new property, `scm.server.proxy.timeout`, has been added for configuring the Descriptor fetch timeout in the Cloudera Manager Admin Console. This is useful when tuning Cloudera Manager for very large deployments. Previously, this value was configured at the service level in various Advanced configuration snippets.

You can find the property by navigating to **Administration > Settings**.

Cloudera Bug: [OPSAPS-41578](#)

• CSD Health Reporting

Added support for determining the health of a CSD service based on the health of its roles. For more information, see the Health Aggregation section in the [CSD Documentation](#).

- **New Validator for Banned YARN Users**

A new validator has been added to ensure that the **Banned System Users** list is the same across all YARN NodeManagers when using Kerberos authentication. The YARN service does not start if the validation fails. To see the list of banned users, select the YARN cluster and navigate to **Configuration**. Search for the `banned.users` property.

- **Resume Rolling Upgrade**

When running a Rolling Restart as part of an upgrade, you can resume the rolling restart after fixing problems that caused the upgrade to fail because one or more hosts did not successfully restart. After you fix the problems you can now resume the rolling restart and Cloudera Manager will skip restarting roles that have already successfully restarted. This change speeds up retrying rolling restarts for large clusters.

- **Diagnostic Bundles Collection for Upgrade Failures**

During a cluster upgrade, if there is a failure, Cloudera Manager now allows you to send a diagnostic bundle to Cloudera support. The Upgrade Wizard opens the **Send Diagnostic Data** dialog box with the current cluster name and time duration pre-populated.

- **New Impala metrics for hedged reads, JVM heap usage and connection setup queue size**

New metrics have been added JVM Heap usage of the Catalog Server and Hedged reads.

- **Support for Sentry with a highly available Hive Metastore**

It is now possible to use HDFS Sentry Sync when running a Hive Metastore using high availability.

- **New placement rules for CSD Services**

A new placement rule for CSD-based services has been added to the Service Descriptor Language called `alwaysWithAny`. When this rule is present, the specified role must always be placed on the same host where the roles specified in the rule are placed. The specified role no longer appears in the wizard when adding this service. Instead, one instance of this role is automatically placed on any host that has at least one of the primary roles. If more than one of the primary roles are themselves placed in the same host, then only one instance of this role is automatically placed on that host. There should be at least two unique primary roles defined in the `alwaysWithAny` rule. And, the `alwaysWithAny` rule is mutually exclusive to the `alwaysWith` rule and they should not be defined together for the same role. If a user assigns roles in a way that violates this placement rule, the service shows a configuration error and fails to start. See the [Service Descriptor Language Reference](#).

- **Navigator search and tagging are now on by default**

The Navigator search and tagging features of Hue are now enabled by default when adding a new Hue service to a cluster running CDH 5.12 or higher.

- **Add `protocol`, `accept_count`, and `acceptor_thread_count` parameters to LUNA_KMS and THALES_KMS CSDs**

New performance tuning parameters related to Tomcat have been exposed within the KMS services designed to work with the Luna and Thales Hardware Security Modules (HSM). These parameters only take effect when running CDH5.12.1 and higher. The following parameters were added:

- `protocol`
- `accept_count`
- `acceptor_thread_count`

What's New in Cloudera Manager 5.12.2

Cloudera Manager 5.12.2 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.12.2](#) on page 2576.

What's New in Cloudera Manager 5.12.1

Cloudera Manager 5.12.1 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.12.1](#) on page 2578.

What's New in Cloudera Manager 5.12.0

- **Backup and Disaster Recovery**

- **Refreshing Impala metadata during replication**

You can now use an option in the Cloudera Manager Admin Console to configure BDR to automatically refresh Impala's metadata cache in the destination cluster during replication. Previously, this feature required an **Advanced Configuration Snippet (Safety Valve)**. See [Invalidating Impala Metadata](#).

- **Automatic renewal of Kerberos tickets and Delegation Tokens**

Previously, BDR replication jobs would fail on a Kerberized cluster if the job duration was longer than the renewal interval for the HDFS delegation token. With this fix, both the delegation token and Kerberos ticket are renewed until the max lifetime of token/ticket (default value is 7 days). This enables longer replications without needing to bring down the source cluster to change the ticket timeout.

- **Streamlined Kerberos Configuration**

- As part of **Test Connectivity** for peers, Cloudera Manager now tests for properly configured Kerberos authentication on the source and destination clusters. **Test Connectivity** runs automatically when you add a peer for replication, or you can manually initiate Test Connectivity from the **Actions** menu. This feature is available when the source and destination clusters run Cloudera Manager 5.12 or higher. See [Enabling Replication Between Clusters with Kerberos Authentication](#).
 - If Cloudera Manager is managing the Kerberos configuration (`krb5.conf`) for your clusters, BDR can automatically make some required changes to your Kerberos configuration based on issues found during the **Test Connectivity** action.
 - The configuration process for adding peers when using Kerberized clusters is simplified if both the source and target clusters use Cloudera Manager 5.12 or later. Now, you only need to setup trust on the target cluster and not the source, reducing the complexity of enabling Hive Replication. See [Enabling Replication Between Clusters with Kerberos Authentication](#)

- **Add a name and description to replication schedules**

When you create or edit a replication schedule, you can add a name on the **General** tab and add a description on the **Advanced** tab.

- **Hive Metastore Schema Integrity Checker**

Cloudera Manager now uses the Hive Metastore schemaTool for validating the integrity of Hive metadata. When you upgrade a cluster that contains a Hive Service to CDH 5.12 or higher using the Cloudera Manager Upgrade Wizard or command line, before upgrading the Hive metastore schema, Cloudera Manager first runs a validation check to detect any corruption. If the validation check fails, Cloudera Manager displays the error and stops the upgrade. Corruption issues should be resolved before proceeding with the upgrade.

- **Support for HSM Key Provider**

The HDFS Encryption Wizard in Cloudera Manager now supports configuration of the Hardware Security Module (HSM) Key Providers supported by CDH 5.12 for encryption key management.

- **Sending Diagnostic Bundles**

The user interface in the Cloudera Manager Admin Console for collecting and sending diagnostic bundles has been improved. Regardless of how diagnostic data collection is configured before you start, each time you create a bundle, you can now select one of the following options: **Collect and Upload Diagnostic Data to Cloudera Support** or **Collect Diagnostics Data only**. Additionally, the Cloudera Manager Admin console better indicates the status of the bundle. For example, showing whether or not the bundle was successfully sent to Cloudera.

- **Delete Kerberos Service Principals**

You can now delete MIT Kerberos or Active Directory Service Principals that were previously created by Cloudera Manager while Kerberizing a cluster using the `delete_credentials` API.

- **HBase Region in Transition Health Check**

Cloudera Manager now performs a health check to detect whether HBase regions have become stuck in transition during splitting and merging operations.

- **Replication factor for MapReduce job submission files**

New auto-configuration logic for MR1 and MR2's **Submit Replication Factor** property attempts to choose a value that is at least the value of the HDFS Replication Factor for clusters with three or more DataNodes. Additionally, a new configuration validator raises a configuration warning if the existing Submit Replication Factor is lower than the HDFS Replication Factor if the cluster has at least 3 DataNodes.

- **Custom Header Color**

You can customize the header color that Cloudera Manager displays in the web UI. Select **Administration > Settings**. Select **Other** for the **Category** and use the drop-down menu for **Custom Header Color**.

- **Dynamic Resource Pools UI**

The **Dynamic Resource Pools** user interface now displays Access Control information about resource pools, showing whether they are freely usable, restricted to a custom set of users/groups, or inherit ACLs from their parent pool.

- **Example Impala Shell Command**

The Impala Service Status Page now includes an example Impala Shell Command.

- **Configurable S3 Endpoint**

The S3 Connector service now allows you to configure the default S3 endpoint used by HDFS clients (including Hive and Impala), ensuring all S3 data created/accessed by your cluster is (by default) stored in the AWS region of your choice. Additionally, Hue is configured to automatically use the default endpoint as the S3 Connector.

- **Solr**

- **Request Rate and Index Size Charts**

The graphs on the Solr status page now include the request rates against the service and the aggregate size of the indices.

- **New Tags in Solr Logs**

The logging for Solr has been improved. Logs now include the following IDs: thread, shard, replica, and collection.

What's New in Cloudera Manager 5.11.2

Cloudera Manager 5.11.2 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.11.2](#) on page 2585.

- **New Tags in Solr Logs**

The logging for Solr has been improved. Logs now include the following IDs: thread, shard, replica, and collection.

What's New in Cloudera Manager 5.11.1

Cloudera Manager 5.11.1 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.11.1](#) on page 2587.

What's New in Cloudera Manager 5.11.0

- **Amazon S3**

- **Amazon S3 Consistency with Metadata Caching (S3Guard)**

Data written to Amazon S3 buckets is subject to the "eventual consistency" guarantee provided by S3, which means that data written to S3 may not be immediately available for queries and listing operations. This can cause failures in multi-step ETL workflows, where data from a previous step is not available to the next step. To mitigate these consistency issues you can now configure metadata caching for data stored in Amazon S3 using S3Guard. Some workloads that access S3 may also see modest performance improvements with metadata caching. S3Guard requires that you provision a DynamoDB database from Amazon Web Services and configure S3Guard using the Cloudera Manager Admin Console or command-line tools. See [Configuring and Managing S3Guard](#).

- **Operating System Support**

- **SLES 12 SP2 Support**

SLES 12, SP2 is now supported as of Cloudera Manager and CDH 5.11 and higher.

- **Mixed Operating system support for gateway hosts running Cloudera Data Science Workbench**

A Gateway host that is dedicated to running Cloudera Data Science Workbench can use RHEL/CentOS 7.2 even if the remaining hosts in your cluster are running any of the other [supported operating systems](#). All hosts must run the same version of the [Oracle JDK](#).

- **Backup and Disaster Recovery**

- **Refreshing Impala metadata during replication**

You can now configure Hive/Impala replication jobs to run the `INVALIDATE METADATA` Impala statement in the destination cluster automatically at the end of the replication process, allowing newly replicated data to be immediately queried by Impala. See [Invalidating Impala Metadata](#).

- **Hive Replication to Amazon S3 now supported for regions that support only Signature Version 4 signing protocol**

Replications from Hive or HBase to Amazon S3 are now supported for S3 regions that only support Amazon's **Signature Version 4** signing protocol. You must add the `fs.s3a.endpoint` property to the **Cluster-wide Advanced Configuration Snippet (Safety Valve) for core-site.xml** property and set its value to the Amazon S3 region. For example:

```
<property>
<name>fs.s3a.endpoint</name>
<value>s3.us-east-2.amazonaws.com</value>
</property>
```

You can access this property in Cloudera Manager at **Home > Configuration > Advanced Configuration Snippets**.

- **Peak Memory Usage Filter now tracked per container for YARN applications**

Peak container memory usage is now tracked for YARN applications and new filter attribute, **Used Memory Max** has been added for monitoring YARN applications.

- **Improved Kerberos-Encryption-Type Handling by Cloudera Manager**

Cloudera Manager validates the Kerberos encryption type as it is being entered into the Cloudera Manager Admin Console, and displays an error message if the type is not a valid MIT or Microsoft Active Directory (Kerberos) encryption type. Administrators can disable the feature when necessary—for example, if new encryption types added to Kerberos are ahead of the encryption types supported by Cloudera Manager (invalid encryption types fail, regardless of warning message display).

- **Enabling SPNEGO authentication for Hue**

Enabling the Hue **Authentication Backend** property (for SPNEGO) now automatically adds all necessary environments and kerberos credentials. Previously, you needed to follow this procedure: [Enabling SPNEGO as an Authentication Backend for Hue](#).

- **New and Changed Configuration**

- **Auto-configuration of HBase Thrift Authentication when Kerberos is enabled**

When Kerberos is enabled for the cluster, the value of the HBase configuration parameter **HBase Thrift Authentication** is automatically set to `auth-conf`. Clusters that already have Kerberos enabled will not have this setting changed when upgrading Cloudera Manager, this only affects enabling Kerberos in the future.

- **New HDFS NameNode configuration property for deleting the trash**

A new HDFS NameNode property, **Filesystem Trash Checkpoint Interval** (`fs.trash.checkpoint.interval`) has been introduced with a default value of 1 hour. This property causes the NameNode better respect and accurately enforce the configured HDFS trash deletion interval set with the **Filesystem Trash Interval** property (`fs.trash.interval`).

The old behaviour without this property accidentally caused many files in the HDFS trash to be deleted only when twice the desired trash deletion interval had transpired because the checkpoint interval matched the deletion interval. If the older implicit behaviour of retaining trash files for a longer time is desired, consider raising the value of the **Filesystem Trash Interval** property to a more suitable value. When upgrading to this version of Cloudera Manager or when changing this property, all HDFS NameNode role instances will be marked stale, and you must restart the HDFS NameNode role instances and their dependent services for this change to take effect.

- **New Auto Logout Timeout property for Hue**

A new configuration property has been added for the Hue service. The **Auto Logout Timeout** property controls how long the Hue browser can remain idle before automatically logging out the user. Set the property to `-1` to disable automatic logout. To configure the property, go to the Hue service, select the **Configuration** tab and search for the property.

- **New performance tuning properties for Key Management Server (KMS)**

The following new properties have been added for tuning the performance of the KMS service:

- **KMS Accept Count**
- **KMS Handler Protocol**
- **KMS Acceptor Thread Count**

- **New API endpoint for refreshing parcel information**

A new REST API endpoint has been added to refresh parcel information from both local and remote repositories. The endpoint URL is:

```
/api/v16/cm/commands/refreshParcelRepos
```

See [Cloudera Manager API documentation](#)

- **New Metrics and Health Tests for Service Monitor and Host Monitor metric collection**

A new metric, **mgmt_aggregation_run_duration**, has been added to the Service Monitor and Host Monitor metrics to indicate how much time it takes to store metrics collected in last minute. This metric can be used to determine if more heap or non-heap memory is needed for these roles.

See [Cloudera Manager Metrics](#).

New Health Tests, **Host Monitor Metrics Aggregation Run Duration Test** and **Service Monitor Metrics Aggregation Run Duration Test** have also been added to detect potential resource configuration issues with service monitor and host monitor.

See [Cloudera Manager Health Tests](#).

- **New validation for YARN NodeManager log directory**

Cloudera Manager now validates whether all YARN NodeManagers are storing logs in the same distributed file system directory so that no logs are missing from Job History Server. If NodeManagers have different configuration values, there will be a configuration error after upgrading Cloudera Manager to 5.11.

- **New Dynamic Resource Pool option**

Configuration of Nested User Pools (except existingSecondaryGroup) now includes a **Create pool if it does not exist** checkbox to indicate whether to create a sub-pool.

What's New in Cloudera Manager 5.10.2

Cloudera Manager 5.10.2 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.10.2](#) on page 2592.

What's New in Cloudera Manager 5.10.1

Cloudera Manager 5.10.1 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.10.1](#) on page 2593.

What's New in Cloudera Manager 5.10.0

- **Backup and Disaster Recovery**

- **Change of default behavior for Impala metadata setting in Hive replications**

There are now three options for configuring replication of Impala metadata for Hive replication jobs:

- **No** – Impala metadata is not replicated.
- **Yes** – Impala metadata is replicated.
- **Auto** – Cloudera Manager decides the value for this option based on the version of CDH in your cluster.

In Cloudera Manager 5.9, when creating a Hive replication schedule, the **Replicate Impala Metadata** option was not selected by default (*false*). In Cloudera Manager 5.10, the value defaults to **Auto**, so that Cloudera Manager decides what the value should be.

- **Replication of Impala Cached Column statistics**

Table and partition-level column statistics stored in the Hive metastore and used by Impala are now replicated during Hive Replication. This is supported between a replication source with Cloudera Manager running version 5.10 or higher and a replication target running Cloudera Manager 5.10 or higher. Because this change replicates more information, the same schedule may take more time to complete if column statistics are present.

- **Performance summaries of HDFS and Hive replication jobs**

You can download full performance reports for HDFS and Hive replications from the **Replication Schedules** page and from the **Replication History** page. You can also filter the output by error, deleted, or skipped status. See [Monitoring the Performance of Hive/Impala Replications](#) and [Monitoring the Performance of HDFS Replications](#).

- **HDFS replication performance monitoring now reports an initial sample**

HDFS performance reports now display early samples during the start of a replication job to give users earlier information about the progress of the job.

- **Scheduler Pool field for replications moved to Resources tab**

When creating replication schedules, the **Scheduler Pool** input field has been moved to the **Resources** tab.

- **Amazon S3**

- **Amazon S3 Object Store**

You can configure auxiliary storage in your cluster using Amazon Simple Storage Service (S3). Client applications such as YARN, MapReduce, or Spark can access data stored in Amazon S3 using Amazon S3 URLs and credentials. See [Configuring the Amazon S3 Connector](#).

- **Amazon S3 Server-side Encryption (SSE-S3)**

Clusters that use Amazon S3 storage can encrypt data using Amazon [server-side encryption \(SSE-S3\)](#). Use Cloudera Manager Admin Console to configure the cluster to use this new feature as detailed in [How to Configure Encryption for Amazon S3](#).

- **Auto-configuration of Hue with AWS credentials (S3 access)**

Cloudera Manager now automatically configures S3 credentials for Hue in the `hue.ini` file.

- **Resource Management**

- **maxResources per user now configurable in YARN Dynamic Resource Pools**

YARN Dynamic Resource Pools now support default capacity limits which automatically apply to all child pools of a resource pool. Using these limits, you can now effectively control the YARN resources available to any user or group in the cluster by configuring these settings on a parent pool and then using placement rules to auto-create child pools per user or group.

- **Key Management Service (KMS)**

- **Java Option configuration added for KMS and Key Trustee KMS**

Additional Java Configuration Options can now be set for the KMS and Key Trustee KMS using Advanced Configuration Snippets. This allows customized Java configuration options for the KMS process.

- **Revised ACLs for KMS**

The ACL set generated by the KMS installation wizard has been updated to implement the recommended secure ACL policies for common key names. For more information, See [Configuring KMS Access Control Lists \(ACLs\)](#).

- **Security**

- **Redaction of stdout and stderr logs**

Redaction logic now applies to the output streams of a service to stdout and stderr logs to redact sensitive information from these files.

- **Encryption of Cloudera Manager database password**

By default the Cloudera Manager database password is stored as clear text in the `/etc/cloudera-scm-server/db.properties` file. You can now specify a program to execute whose standard output is the password. If `com.cloudera.cmf.db.password` is not found in `/etc/cloudera-scm-server/db.properties`, then the property `com.cloudera.cmf.db.password_script`, if it exists, is used. The value of this property is executed as a program, and the value returned to `stdout` is used as the password.

- **Monitoring**

- **Collection of Metrics can be disabled for specified roles**

A new configuration has been introduced at the role level to disable monitoring by the Cloudera Manager Agent for each individual role. By default, monitoring is enabled for all roles. Once disabled, you must restart the role to make the change take effect. This can help in scenarios where the Cloudera Manager Service Monitor is experiencing performance issues and does not have enough resources. See [Disabling Metrics for Specific Roles](#).

- **New chart display number of monitored entities**

A new chart displays on the Cloudera Management Service status page that shows the number of monitored entities by service monitor and host monitor.

- **Default timeouts configurable for Service Monitor, Host Monitor, Reports Manager, Events Server, and Activity Monitor**

The default timeout is now configurable for Service Monitor, Host Monitor, Reports Manager, Events Server, and Activity Monitor to register themselves to a cluster. Users can change the timeout by increasing the values of the **Descriptor Fetch Max Tries** property (the default value is 5) and the **Descriptor Fetch Tries Interval** property (the default value is 2).

- **New Health Test for Service Monitor heap usage.**

A health test has been added to monitor heap usage of the Service Monitor. This test can be useful when diagnosing out of memory errors in the Service Monitor.

- **Users can send feedback on Cloudera Manager**

You can go to **Support > Feedback** to send feedback to Cloudera about Cloudera Manager.

- **Suppression of warning about embedded database for Cloudera Manager.**

When Cloudera Manager is configured to use the embedded PostgreSQL data base, it displays a banner warning that the embedded PostgreSQL database is not recommended for production environments. You can now suppress this banner: Go to **Administration > Settings > Enable Embedded Database Check**.

- **New HDFS balancer related configuration options**

HDFS balancer can now be configured to specify which hosts are included and excluded or which hosts are used as sources for transferring replicas. Additional properties for tuning the performance of the balancer can now also be configured starting for CDH 5.10.0 and higher.

The following configuration will be marked stale: `dfs.datanode.balance.max.concurrent.moves`. You can safely ignore the warning and defer restarting.

- **Removing service dependencies**

Previously, when the user tries to delete a service that another service depends on, Cloudera Manager displays a dialog box explaining that the dependent service should be deleted. There is now a link that takes user to the configuration page where they can remove this dependency.

- **HBase topology files now deployed in client configurations**

When the HBase property **Enable Replication To Secondary Region Replicas** is enabled, the `topology.py` and `topology.mapp` files are deployed with the client configuration and the **Topology Script File Name** property is set to the deployed `topology.py` path.

- **The length of Impala queries retained by the Service Monitor is now configurable.**

The maximum size of impala queries is now configurable and the limit the default size to 10k chars.

What's New in Cloudera Manager 5.9.3

Cloudera Manager 5.9.3 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.9.3](#) on page 2597.

What's New in Cloudera Manager 5.9.2

Cloudera Manager 5.9.2 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.9.2](#) on page 2600.

What's New in Cloudera Manager 5.9.1

Cloudera Manager 5.9.1 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.9.1](#) on page 2601.

What's New in Cloudera Manager 5.9.0

- **Creating Virtual Machine Images**

Documentation has been added with procedures to create virtual images of Cloudera Manager and cluster hosts. See [Creating Virtual Images of Cluster Hosts](#).

- **Security**

- **External/Cloud account configuration in Cloudera Manager**

Account configuration for access to Amazon Web Services is now available through the centralized UI menu **External Accounts**.

- **Key Trustee Server rolling restart**

Key Trustee Server now supports rolling restart.

- **Backup and Disaster Recovery**

- You can now replicate HDFS files and Hive data to and from an Amazon S3 instance. See [HDFS Replication To and From Cloud Storage](#) and [Hive/Impala Replication To and From Cloud Storage](#).

- There are some new tuning options to improve performance of HDFS replication. See [HDFS Replication Tuning](#).

- You can now download performance data about HDFS replication jobs from the **Replication Schedules** and **Replication History** pages. See [Monitoring the Performance of HDFS Replications](#).

- Hive replication now stores Hive UDFs in the Hive metastore. [Replication of Impala and Hive User Defined Functions \(UDFs\)](#).

- The user interface for creating replication schedules has been reorganized to present the configuration options on three tabs: **General**, **Resources**, and **Advanced**.

- **Uncheck Replicate Impala Metadata by default**

When creating a Hive replication schedule, the option **Replicate Impala Metadata** was checked (*true*) by default. In Cloudera Manager 5.9 and higher, the value is unchecked (*false*) by default.

- **YARN BDR enhancement**

YARN jobs now include the BDR schedule ID that launched the job so you can connect logs with existing schedules, if multiple schedules exist.

- **Resource Management**

- **Custom Cluster Utilization Reports**

Documentation has been added to create custom Cluster Utilization reports that you can export data from. See [Creating a Custom Cluster Utilization Report](#).

- **New settings for continuous scheduling**

For new installs, default values for configurations have been changed.

`yarn_scheduler_fair_continuous_scheduling_enabled` is set to *false*.

`resourcemanager_fair_scheduler_assign_multiple` is set to *'true'*. Existing settings are preserved when you upgrade from a lower version.

- **YARN historical reports by user show pool-user entity**

When Cloudera Manager manages multiple clusters, there is no per user tracking for historical applications and queries across clusters. Instead, **Historical Applications by User** and **Historical Queries by User** show applications and queries per user and pool. (A pool is associated with a specific cluster.)

- **Directory Usage Report needs export capability**

Directory usage reports can be exported as a CSV file.

- **Cloudera Manager Admin Console User Interface**

- **Service colors**

A new set of colors is used to represent each kind of service.
- **Move the table sorting icon to the right**

The table sorting icon now appears consistently on the right hand side of each column.
- **Improved Configuration Diff Display**

Changes displayed in the configuration history page are much more user friendly. For a large section of changed text, Cloudera Manager generates a diff between the old and the new and displays the diff.

When a user changes only the password, Cloudera Manager does not show the delta: both the old and the new passwords are masked out before the comparison is performed.
- **Move actions menu to the top header**

The actions menu now appears next to the entity title.
- **Move Federation and High Availability to a separate page**

The **Federation** and **High Availability** sections used to appear on the **HDFS Instances** page of an HDFS service. They have been moved to a new page called **Federation and High Availability**. There is a link from the existing **Instances** page to this new page.
- **Remove repeated heading below the second level navigation**

Subtitles below the second level navigation tabs are removed because they repeated the content in the tabs.
- **Move maintenance mode and badges to the title area**

Maintenance mode, staleness badges now appear next to the title of the entity.
- **Express wizard allows you to add Kafka**

Kafka is now listed in the custom services when you click the **Add Cluster** button.
- **Cloudera Manager API**
 - **Add update_user to Python API client**

Added the `update_user()` method to the Python API client `api_client.py`.
 - **Expose API endpoint to add a specific path**

New API endpoints have been added that allow users to add, list and remove Watched Directories in HDFS service.
- **Logging**
 - **Include host in log file name**

Kafka log4j log files now include the host name in the format `kafka-broker- $\{host\}$.log`. Similarly, MirrorMaker logs now include the host name in the format `kafka-mirrormaker- $\{host\}$.log`. Due to the log file name change, when you upgrade Cloudera Manager it no longer recognizes your old log files in log search, though they are still present on disk.
 - **Configuration changes to Cloudera Manager audit log**

Cloudera Manager displays the **History** and **Rollback** support for the Cloudera Manager Settings. (**Administration > Settings**). This helps you to track the changes made by an administrator so that Cloudera Support can provide better service when certain Cloudera Manager administrative settings are modified.
- **Diagnostic Bundles**
 - **Show the Diagnostic Bundle Redaction Policy using the redaction config**

You can specify what information should be redacted in the diagnostic bundle in the UI using **Administration > Settings > Redaction Parameters for Diagnostic Bundles**.

- **Upgrade**

- **Report that a simple restart was performed if rolling restart could not be performed**

Informs you when a simple restart is performed instead of rolling restart on a service because rolling restart is not available.

- **Oozie**

- **Provide dump / load functionality for Oozie DB**

The **Actions** menu in the Oozie service has two new commands, **Dump Database** and **Load Database**. These commands make it easier to migrate an Oozie database to another database supported by Oozie. The **Dump Database** command exports Oozie's database to a file (configurable by **Database Dump File** setting). **Load Database** loads the file into a database.

- **Install Oozie ShareLib permissions change**

Install Oozie ShareLib Command assigns correct permissions to the uploaded libraries. This prevents breaking Oozie workflows with a custom umask setting.

- **Configuration Changes**

- **Solr zkClientTimeout option**

Added the *zkClientTimeout* parameter for ZooKeeper.

- **Add JHIST compression as a configuration option**

Added a new option for setting the file format used by an ApplicationMaster when generating the .jhist file.

- **Enable heap dump by default for all daemons**

Starting in version 5.9, when you configure roles that are JVM based, the *Dump Heap When Out of Memory* configuration parameter defaults to *true*. An upgrade from a pre-5.9 version maintains your pre-5.9 settings.

- **Cloudera Manager support for client-side YARN graceful decommissioning**

Adds the ability to perform a graceful decommission on YARN NodeManager roles whereby the Node Manager is not assigned new containers, and waits for any currently running applications to finish before being decommissioned unless a timeout occurs. You can configure the timeout using the *Node Manager Graceful Decommission Timeout* configuration property in the YARN Service. The default behavior has not changed, and continues to be a non-graceful decommission. Affects Cloudera Manager 5.9.0 and higher, and CDH 5.9.0 and higher.

- **Deploy Client Configuration command details page now shows stdout/stderr**

stdout and *stderr* log links are now shown in the UI when there is a failure while deploying client configurations.

- **Make EXTRA_RATIO configurable for Headlamp indexing**

Added the configuration parameter, *Extra Space Ratio for Indexing*, to Reports Manager. You can use the parameter to make the speed of indexing faster by allocating additional memory.

- **Configure HBase Indexer to wait longer for ZooKeeper to come up**

The default amount of time that HBase Indexer roles attempts to connect to ZooKeeper has been increased from 30 to 60 seconds. This default can be adjusted by setting a new Cloudera Manager configuration parameter, HBase Indexer ZooKeeper Session Timeout.

- **Embedded database mode improvements**

In version 5.9 and higher, Cloudera Manager can clearly identify whether or not a customer is using the embedded PostgreSQL database. Cloudera does not recommend the embedded database for production use, and requests that customers deploy production systems using an external database. The diagnostic bundles now contain information about whether or not a customer is using the embedded PostgreSQL database. Support can then reach out to customers accordingly.

If Cloudera Manager is configured to use the embedded PostgreSQL database, a yellow banner appears in the UI recommending that you upgrade to a supported external database.

- **Fix CatalogServiceClient to handle TLS connections to catalogd for UDF replication**

When Impala uses SSL, we now support TLS Connection to Catalog Server. Customers can enable replication for any Impala UDFs/Metadata (in Hive Replication) in Cloudera Manager 5.9 and higher.

- **Do not show steps that are unreachable (skipped)**

When running wizards from the Cloudera Manager Admin Console that add a cluster, add a service, perform an upgrade, and other tasks, steps do not display when they are not reachable or do not apply to the current configuration.

- **Improve Cloudera Manager provisioning performance on AWS**

Add support for resetting Cloudera Manager GUID/UUID. This is accomplished by checking the UUID file.

If Cloudera Manager finds the UUID file (`/etc/cloudera-scm-server/uuid`) and the UUID is different than the GUID in the `cm_version` table, it updates the GUID in the `cm_version` table with the contents of the UUID file and removes the UUID file.

What's New in Cloudera Manager 5.8.5

Cloudera Manager 5.8.5 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.8.5](#) on page 2606.

What's New in Cloudera Manager 5.8.4

Cloudera Manager 5.8.4 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.8.4](#) on page 2608.

What's New in Cloudera Manager 5.8.3

Cloudera Manager 5.8.3 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.8.3](#) on page 2609.

What's New in Cloudera Manager 5.8.2

Cloudera Manager 5.8.2 is a maintenance release with many fixed issues. See [Issues Fixed in Cloudera Manager 5.8.2](#) on page 2610.

What's New in Cloudera Manager 5.8.1

An issue has been fixed. See [Issues Fixed in Cloudera Manager 5.8.1](#) on page 2612.



Note: Although there is a Cloudera Manager 5.8.1 release, there is no synchronous CDH 5.8.1 release.

What's New in Cloudera Manager 5.8.0

- **Operating Systems** - Support for Debian 8.2.
- **Resource management and utilization** - Added support for nesting dynamic resource pools within a named pool at runtime.
- **Backup and Disaster Recovery**

- The **Replication Schedules** page now has a search function for finding scheduled replications.
- You can now specify a start and end time for the events that are included in manually-triggered diagnostic bundles. See [Manually Triggering Collection and Transfer of Diagnostic Data to Cloudera](#).
- **Impala**
 - Impala adds a new configuration option, `Use HDFS Rules to Map Kerberos Principals to Short Names`. Enabling this option makes Impala pickup `hadoop.security.auth_to_local` configuration from HDFS configurations and uses it for Kerberos principal-to-short-name translation. This only applies for Cloudera Manager 5.8.0 and higher and CDH 5.8.0 and higher. It only affects deployments where Impala is set up to use Kerberos as the authentication mechanism. It defaults to `false`, to preserve the behavior from earlier CDH versions. This has no impact on upgrade.
 - Enable Impala Admission Control and Enable Dynamic Resource Pools are now enabled by default. Customized configuration values are preserved during the upgrade.
 - Impala Admission Control now supports a global method for editing the Access Control List.
- **EMC Isilon**
 - Kerberos is now fully supported for replications between clusters using Isilon storage. You must configure a custom principal.
- **Security**
 - **Active Directory KDC**
 - Active Directory account properties are now configurable from Cloudera Manager's **Administration > Settings** page.
 - It is now possible to use Cloudera Manager to regenerate principals for clusters using an Active Directory KDC. Cloudera Manager 5.8 includes a new configuration called **Active Directory Delete Accounts on Credential Regeneration**. Enabling this will allow Cloudera Manager to automatically delete existing AD accounts and complete the regeneration process.
 - Cloudera Manager now allows you to configure the encryption types (or `encryptype`) used by an Active Directory KDC to protect its data.
 - **Redaction**: In the Cloudera Manager Admin Console, Advanced Configuration Snippet parameters will now be redacted to block sensitive information such as passwords or secret keys.
 - **Sentry**
 - Cloudera Search adds support for storing permissions in the Sentry service. You can enable storing permissions in the Sentry service by [Enabling the Sentry Service for Solr](#). If you have already configured Sentry's policy file-based approach, you can migrate existing authorization settings as described in [Migrating from Sentry Policy Files to the Sentry Service](#). `solrctl` has been extended to support:
 - Migrating existing policy files to the Sentry service
 - Managing managing permissions in the Sentry service
 - Sentry supports data stored on Amazon S3 and can secure URIs with an S3 schema.
- **YARN**
 - YARN Allowed System Users now includes `hbase` by default. This is helpful when running certain tools for HBase that need to execute MapReduce jobs.

What's New in Cloudera Manager 5.7.6

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.7.6](#) on page 2621.

What's New in Cloudera Manager 5.7.2

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.7.2](#) on page 2622.

What's New in Cloudera Manager 5.7.1

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.7.1](#) on page 2624.

What's New in Cloudera Manager 5.7.0

- **Operating Systems** - Support for:

- RHEL/CentOS 6.6, 6.7, 7.1, and 7.2
- Oracle Enterprise Linux (OEL) 7.1 and 7.2
- SUSE Linux Enterprise Server (SLES) 11 with Service Packs 2, 3, 4
- Debian: Wheezy 7.0, 7.1, and 7.8



Important: Cloudera supports RHEL 7 with the following limitations:

- Only RHEL 7.3, 7.2 and 7.1 are supported. RHEL 7.0 is not supported.
- RHEL 7.1 is only supported with CDH 5.5 and higher.
- RHEL 7.2 is only supported with CDH 5.7 and higher.
- Only new installations of RHEL 7.2 and 7.1 are supported by Cloudera. For upgrades to RHEL 7.1 or 7.2, contact your OS vendor and see [Does Red Hat support upgrades between major versions of Red Hat Enterprise Linux?](#)

- **Resource management and utilization**

- Simplified and expanded resource management. The screens for YARN and Impala dynamic resource pools are now managed separately.
- Resource pools now support the `allowPreemptionFrom`, `minSharePreemptionTimeout`, and `fairSharePreemptionTimeout` attributes.
- Cluster utilization reports track usage of resources allocated using dynamic resource pools.
- The new **Directory Usage Report** now shows aggregated usage information, including quotas and file sizes, which are sortable. You can also perform multiple actions on filesystem objects. .
- Two new predicates have been added to the tsquery language: `day in` and `hour in`. These allow you to limit streams to specified days of the week and specified hours of each day, respectively.

- **Extensibility** - For more information, see [Cloudera Manager Extensions](#).

- Parcels are typed according to the OS version. The parcel extension indicates the version. The library for developing external parcels now supports an extension for RHEL 7.
- A new environment variable, `ZK_PRINCIPAL_NAME`, is now defined for CSD processes when ZooKeeper is Kerberized and has a custom principal.
- A new flag, `jvmBased`, is now available to CSD authors to indicate that a CSD role is JVM-based. This flag enables a set of JVM-related features in Cloudera Manager—for example, the ability to automatically generate a heap dump when an Out Of Memory error occurs.

- **API**

- Cloudera Manager now attempts to gracefully handle overlapping API calls.
- All distributed filesystem services (such as HDFS or Isilon) installed in a cluster can now be enumerated using the API.
- You can export the complete configuration of a CDH cluster managed by Cloudera Manager as a template, modify the template, and import the template to create a new cluster.

- The advanced configuration snippet editor now allows you to edit properties as name/value pairs. This is the default, however you can also choose to edit the snippet as XML.
- HBase now includes metrics and charts for replication. These charts are available in the Chart Library for each RegionServer.

- When you click the **Role Log** link when viewing a role, the log is opened at the current timestamp, rather than the top of the log file. This enables you to see the relevant log messages when investigating an event that occurred recently.
- A new `tsquery` function called `counter_delta` has been added to accurately compute the difference between consecutive data points for counter metrics.
- The `distcp` utility now supports setting records per chunk, using the `distcp.dynamic.recordsPerChunk` in an advanced configuration snippet to set the number of records (paths) in each chunk. When a value is set for `distcp.dynamic.recordsPerChunk`, other related settings, such as the maximum number of chunks tolerable, the ideal number of chunks, and the split ratio, are ignored.
- A warning is shown when upgrading with incompatible versions of Kafka and CDH. The Kafka client libraries bundled in CDH cannot communicate with an older Kafka server.
- You can override the `sudo` commands that Cloudera Manager agent uses by redirecting the `sudo` commands to a script that you write to allow or disallow certain actions.
- **Hive**
 - Hive on Spark is now supported.
 - The default execution engine for Hive can now be configured, which makes it easy to run all Hive jobs on Spark.
 - HiveServer2 now has a Web UI.
 - Hive and HDFS replication source/target listings now work with Isilon.
 - A dialog box now displays when scheduling Hive replications, reminding you to take snapshots of the Hive warehouse directory.
 - The Direct SQL option is now enabled in Hive Metastore.
 - When upgrading to CDH 5.7, if Hive is configured to use YARN, all Hive on Spark parameters are automatically tuned to recommended values. If Hive on Spark was previously tuned, this is skipped.

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.7.0](#) on page 2627.

What's New in Cloudera Manager 5.6.1

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.6.1](#) on page 2630.

What's New in Cloudera Manager 5.6.0

What's New in Cloudera Manager 5.5.4

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.5.4](#) on page 2632.

What's New in Cloudera Manager 5.5.3

An issue has been fixed. See [Issues Fixed in Cloudera Manager 5.5.3](#) on page 2633.

What's New in Cloudera Manager 5.5.2



Note: If you are upgrading from a previous version of Cloudera Manager, Cloudera recommends that you upgrade to version 5.5.3 or higher.

- New Impala flags added for web server certificate files and passwords. This adds support for the `--webserver_private_key_file` and `--webserver_private_key_password_cmd` flags for the Impala Daemon, the Impala Catalog Server, and the Impala StateStore roles.

A number of issues have also been fixed. See [Issues Fixed in Cloudera Manager 5.5.2](#) on page 2634.

What's New in Cloudera Manager 5.5.1

An issue has been fixed. See [Issues Fixed in Cloudera Manager 5.5.1](#) on page 2637.

What's New in Cloudera Manager 5.5.0

- **Operating Systems** - Support for RHEL/CentOS 6.6 (in SE Linux mode), 6.7, and 7.1, and Oracle Enterprise Linux 7.1.



Important: Cloudera supports RHEL 7 with the following limitations:

- Only RHEL 7.1 is supported. RHEL 7.0 is not supported.
- Only new installations of RHEL 7.1 are supported by Cloudera. For upgrades to RHEL 7.1, contact your OS vendor and see [Does Red Hat support upgrades between major versions of Red Hat Enterprise Linux?](#)

- **Backup and Disaster Recovery (BDR)** - You can use the BDR UI or API to replicate data between encrypted clusters
- **Databases** - Supports MariaDB 5.5, Oracle 12c, and PostgreSQL 9.4.
- Selective service restart after activating parcels is supported.
- Retrying upgrade actions is supported. If a cluster upgrade command fails while in progress, you can retry a command after fixing the cause of failure. On retry, the command restarts from the command step where it failed.
- The command details page for running and recent commands has been redesigned for usability and scalability.
- Instead of serially starting all services for the first time, services that are not dependent are started in parallel. This decreases the time required to start services for the first time after creating a cluster.
- Performance has improved for service startup, client configuration deployment, and calculation of stale configurations.
- **Suppression of notifications**
 - You can suppress the warnings that Cloudera Manager issues when a configuration value is outside the recommended range or is invalid.
 - You can suppress health test warnings.

Suppression can be useful if a warning does not apply to your deployment and you no longer want to see the notification. Suppressed warnings are still retained by Cloudera Manager, and you can unsuppress the warnings at any time.

- **Multi Cloudera Manager Dashboard** - A special mode of Cloudera Manager that enables you to view monitoring data aggregated from multiple Cloudera Manager instances that manage one or more CDH clusters.
- You can decommission roles when services are completely stopped. This allows you to decommission hosts during cluster downtime.
- You can disable collection of certain domain metrics—for example, for HBase RegionServers, Kafka Brokers, and others—through new settings in the host advanced configuration snippet. This is useful in certain support situations and should only be done under the direction of Cloudera Support.
- You can configure which aggregate metrics are automatically generated. This advanced feature can be useful in certain situations to impact the monitoring workload, allowing unused or less-important aggregate metrics to be skipped. This may result in improved performance and the ability to handle larger monitoring workloads, or to retain data for a larger workload for longer. Cloudera recommends using this only under the direction of Cloudera Support.
- Alert Publisher can be configured to pass alert events to a user-defined script. Use this for integrating with other alerting systems or for custom logic (for example, to send some alerts to some people and others to other people).
- Agent minor version mismatches (5.4 to 5.5) now cause bad host health. Maintenance version mismatches (for example, 5.4.x to 5.4.y) still cause concerning host health.
- Cloudera Manager indicates if the Java version in use is too old.
- Cloudera Manager indicates if the supervisor component of the Agent needs to be restarted after an upgrade.
- Full and User Administrators can view active user sessions.
- Full Administrators and Auditors can audit failed and successful logins.
- Multiple user session logins can be disallowed.
- You can configure external authentication so that local administrator emergency access is disabled. This means that no local accounts can log in under any circumstances, including when the external system is not functioning.

- You can turn on authentication for the URLs for downloading client configuration zip files. Previously, authentication was never required.
- Passwords are no longer accessible in cleartext through the Cloudera Manager UI or in the configuration files stored on disk. There are some exceptions; see [Known Issues and Workarounds in Cloudera Manager 5](#) on page 2518.
- **HBase**
 - Use a configuration option in HBase to skip region reload during rolling restart and rolling upgrade, to increase the speed of the operations.
 - HBase rolling restart performance can be improved by increasing the number of Region Mover Threads. If the value of this property is 1, it can lower rolling restart speed. The Admin Console now displays this information and, if the value is 1, advises increasing it.
 - HBase Thrift Server and Rest Server support TLS/SSL.
- **HDFS**
 - HDFS encryption can be enabled using a wizard.
 - Exposes AES as an encryption option for HDFS RPC encryption.
- **Hive**
 - Hive can use TLS/SSL and Kerberos at the same time.
 - When Hive is configured to use TLS/SSL, Hue is automatically configured to use that protocol when communicating with Hive. Similarly, when Impala is configured to use TLS/SSL, Hue is automatically configured to use that protocol when communicating with Impala.
 - HiveServer2 supports a timeout value for idle sessions and operations. By default, it times out client sessions after a week and idle operations after three days. This helps alleviate problems with long-running sessions when using Hue.
 - Cloudera Manager collects and displays various operational metrics for Hive.
- **Hue**
 - Hue supports a Load Balancer role using HTTPD as a load balancer.
 - You can configure certificates trusted by Hue using the TLS/SSL Truststore configuration. This replaces the `REQUESTS_CA_BUNDLE` advanced configuration snippet entry.
 - You can specify a password that protects the Hue private key file.
 - Cloudera Manager collects and displays various operational metrics for Hue. New health tests have been added for Hue as well.
- **Impala** supports TLS/SSL internally between the StateStore and the Catalog Server roles as well as Impala Daemon.
- **Kafka**
 - Kafka supports rolling restart.
 - Kafka displays additional broker metrics.
 - Kafka exposes additional commonly configured parameters.
 - Existing Kafka parameter definitions have updated descriptions, default values, and validation settings.
 - The Kafka broker instance list now shows which broker is the active controller.
- **Key Trustee**
 - The Key Trustee Server CSD is included in Cloudera Manager. Manual installation of the Key Trustee Server CSD is not required.
 - A Key Administrator role in Cloudera Manager is used for configuring HDFS Data at Rest Encryption. Only a Key Administrator and a Full Administrator can make configuration changes to Java Keystore KMS, Key Trustee KMS, and Key Trustee Server. Configuring HDFS to use Data at Rest Encryption is also limited to the Key Administrator and Full Administrator roles. This allows organizations to keep Key Administrators and Cluster Administrators separate, which is a security best practice.
 - When running Key Trustee KMS in a highly available configuration, Cloudera Manager can automatically generate the load balancer URL.

- **Sentry**
 - Sentry introduces column-level access control for tables in Hive and Impala. Previously, Sentry supported privilege granularity only at the table level. You can now assign the `SELECT` privilege on a subset of columns in a table.
 - Sentry supports Kerberos authentication for the Sentry web server.
- **Solr**
 - Solr can be configured with a load balancer in a secure environment.
 - There is a new Solr Max Connector Threads property for Solr Server in CDH 5.1.0 and higher.
 - Solr supports LDAP/AD authentication.
- **Backup and Disaster Recovery**
 - The user interface for scheduling and reviewing replications and snapshots has been improved. You can now view the history of replication jobs and subtasks more easily.
 - When specifying an HDFS replication job, you can apply exclusion filters to exclude specific files or directories.
 - You can download or send to Cloudera Support a diagnostic bundle to troubleshoot replication jobs. Bundles include logs of the replication run. .
 - The performance of the file-listing phase of a replication job has been improved.
 - The performance of the initialization and running phase has been improved.
 - The following advanced configuration snippets for configuring replications have been added:
 - HDFS Replication Advanced Configuration Snippet (Safety Valve) for `hadoop-env.sh`
 - Hive Replication Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`
 - HDFS Replication Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml`
 - HDFS Replication Advanced Configuration Snippet (Safety Valve) for `mapred-site.xml`
 - Snapshot properties for HBase such as thread pool size can be configured in the **HBase Client Advanced Configuration Snippet (Safety Valve) for `hbase-site.xml`** property.
 - Hive partitions are chunked during export and import to avoid message size limitations.
 - Hive replications validate metadata on the destination Hive Metastore before copying HDFS data from the source to avoid copying errors during replication.
 - The use of snapshots to improve replications is documented.
 - The effect of network latency on replications is documented.
 - Scheduled snapshots can be disabled and re-enabled.
 - API improvements:
 - Explicit support for pausing snapshot policies
 - Failed file listing
 - Collection of diagnostic bundles for replication schedules and history

What's New in Cloudera Manager 5.4.10

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.4.10](#) on page 2642.

What's New in Cloudera Manager 5.4.9

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.4.9](#) on page 2643.

What's New in Cloudera Manager 5.4.8

New ability to decommission hosts with stopped services

Adds ability to decommission roles when services are completely stopped. This allows users to decommission hosts during cluster downtime.

A number of issues have also been fixed. See [Issues Fixed in Cloudera Manager 5.4.8](#) on page 2644.

What's New in Cloudera Manager 5.4.7

New service-level advanced configuration snippets for Solr

The following new properties were added:

- Solr Service Advanced Configuration Snippet (Safety Valve) for `core-site.xml`
- Solr Service Advanced Configuration Snippet (Safety Valve) for `hdfs-site.xml`

A number of issues have also been fixed. See [Issues Fixed in Cloudera Manager 5.4.7](#) on page 2644.

What's New in Cloudera Manager 5.4.6

An issue has been fixed. See [Issues Fixed in Cloudera Manager 5.4.6](#) on page 2645.



Note: Although there is a Cloudera Manager 5.4.6 release, there is no synchronous CDH 5.4.6 release.

What's New in Cloudera Manager 5.4.5

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.4.5](#) on page 2646.



Note: Although there is a CDH 5.4.4 release, there is no synchronous Cloudera Manager 5.4.4 release.

What's New in Cloudera Manager 5.4.3

Rollback for CDH 4 to CDH 5 Upgrades

- Rolling back a CDH 4 to CDH 5 upgrade is now supported using Cloudera Manager.

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.4.3](#) on page 2649.



Note: Although there is a CDH 5.4.2 release, there is no synchronous Cloudera Manager 5.4.2 release.

What's New in Cloudera Manager 5.4.1

Hue HA Improvements

- The Cloudera Manager Express and Add Service wizards allow you to add a Hue service with multiple Hue Server roles. For Kerberized clusters, the Add Service wizard automatically adds a collocated Kerberos Ticket Renewer role for each Hue Server role instance.
- When Kerberos is enabled, Cloudera Manager now checks to ensure each Hue Server role is collocated with a Kerberos Ticket Renewer role. If you forget to add a Kerberos Ticket Renewer role when adding a new Hue Server role, a configuration error is generated.

Cloudera Manager High Availability

- High availability for Cloudera Manager is now supported for 5.4.

A number of issues have also been fixed. See [Issues Fixed in Cloudera Manager 5.4.1](#) on page 2655.

What's New in Cloudera Manager 5.4.0

- **OS** - Added support for RHEL 6.6 and CentOS 6.6.
- Cloudera Manager prevents installing or upgrading to a CDH version that is too new for the Cloudera Manager version. When using parcels, it prevents parcel installation. When using packages, it prevents creating services.
- Installation and add service wizards now support the Oozie database.
- New wizard for NameNode, Failover Controller, and JournalNode role migration.

- Parcel page layout redesigned in terms of layout, performance and ease of use. A new parcel per host detail view is added.
- **Configuration**
 - Configuration pages use the new layout by default. The new layout is dramatically improved in terms of layout, performance, and ease of use. The existing layout is accessible via the **Switch to the classic layout** link.
 - New configuration actions:
 - Configuration can now be applied to all clusters as well as for a specific cluster.
 - Several new configuration views have been added to show all non-default values across all clusters and the Cloudera Management Service, as well as differences across all clusters and multiple services of the same type.
 - One-click differences in configuration settings for a specific service across multiple clusters.
- **Support**
 - Include a Cloudera support ticket with YARN application support bundles.
 - Reduce the size of support bundles by specifying log data of interest to include in the bundle.
- **HDFS**
 - Support for HDFS DataNode hot swap.
 - Option to include replication of extended attributes during HDFS replication. HDFS ACLs will now be replicated along with permissions.
- Added support for Hive on Spark.



Important: Hive on Spark is included in CDH 5.4 and higher but is not currently supported nor recommended for production use. To try this feature, use it in a test environment until Cloudera resolves currently existing issues and limitations to make it ready for production use.

- **Security**
 - Secure impersonation support for the Hue HBase app.
 - Redaction of sensitive data in log files and in SQL query history.
 - Support for custom Kerberos principals.
 - Added commands for regenerating Kerberos keytabs at service and host levels. These commands will clear existing keytabs from affected role instances and then trigger the **Generate Credentials** command to create new keytabs.
 - Kerberos support for Sqoop 2.
 - Kerberos and TLS/SSL support for Flume Thrift Source and Sink.
 - Solr TLS/SSL support.
 - Navigator Key Trustee Server can be installed and monitored by Cloudera Manager.
 - HBase Indexer integration with Sentry (File-based) for authorization.

What's New in Cloudera Manager 5.3.10

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.3.10](#) on page 2658.

What's New in Cloudera Manager 5.3.9

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.3.9](#) on page 2658.

What's New in Cloudera Manager 5.3.8

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.3.8](#) on page 2659.

Cloudera Manager 5 Release Notes

What's New In Cloudera Manager 5.3.7

An issue has been fixed. See [Issues Fixed in Cloudera Manager 5.3.7](#) on page 2659.



Note: Although there is a Cloudera Manager 5.3.7 release, there is no synchronous CDH 5.3.7 release.

What's New in Cloudera Manager 5.3.6

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.3.6](#) on page 2660.



Note: Although there is a CDH 5.3.5 release, there is no synchronous Cloudera Manager 5.3.5 release.

What's New in Cloudera Manager 5.3.4

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.3.4](#) on page 2660.

What's New in Cloudera Manager 5.3.3

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.3.3](#) on page 2661.

What's New in Cloudera Manager 5.3.2

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.3.2](#) on page 2662.

What's New in Cloudera Manager 5.3.1

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.3.1](#) on page 2664.

What's New in Cloudera Manager 5.3.0

- **JDK 1.8** - Cloudera Manager adds support for Oracle JDK 1.8.
- **Single user mode** - The Cloudera Manager Agent and all service processes can now be run as a single configured user in environments where running as root is not permitted.
- **CDH upgrade wizard enhanced** - The CDH upgrade wizard now supports minor and maintenance version upgrade as well as major version upgrade.
- **Oozie Sharelib** - The Oozie Sharelib can be updated without restarting the Oozie service.
- **Read-only users prevented from viewing process logs or environment** - Read-only users can no longer view the environment or logs of a process. This is to prevent read-only users from seeing potentially sensitive information.
- New icons for the KMS and Key Trustee services.
- **Data-at-rest encryption**



Important: Cloudera provides two solutions:

- **Navigator Encrypt** is production ready and available to Cloudera customers licensed for Cloudera Navigator. Navigator Encrypt operates at the Linux volume level, so it can encrypt cluster data inside and outside HDFS. Consult your Cloudera account team for more information.
- **HDFS Encryption** is production ready and operates at the HDFS directory level, enabling encryption to be applied only to HDFS folders where needed.

HDFS encryption implements transparent, end-to-end encryption of data read from and written to HDFS by creating encryption zones. An encryption zone is a directory in HDFS with every file and subdirectory in it encrypted. Use *one* of the following services to store, manage, and access encryption zone keys:

- **KMS (File)** - The Hadoop Key Management Server with a file-based Java keystore; maintains a single copy of keys, using simple password-based protection.

- **KMS (Navigator Key Trustee)** - An enterprise-grade key management service that replaces the file-based Java keystore and leverages the advanced key-management capabilities of Cloudera Navigator Key Trustee. Navigator Key Trustee is designed for secure, authenticated administration and cryptographically strong storage of keys on multiple redundant servers that can be located outside the cluster.
- The Cloudera Manager Server now reports the correct number of physical cores and hyper-threading cores if hyper-threading is enabled.
- **Client configurations** - Client configurations are now managed so that they are redeployed when a machine is re-imaged.



Important: The changes to client configurations affect some API calls, as follows:

- When a host ceases to have a client configuration assigned to it, Cloudera Manager will remove it, rather than leaving it behind. If a host has a client configuration assigned and the client configuration is missing, Cloudera Manager will recreate it.
- If you currently use the API command `deployClientConfig` to deploy the client configurations for a particular service, and you pass a specific set of role names to this call to narrow the set of hosts that receive the new client configuration, then you should be aware that:
 - The API command will continue to generate and deploy the client configuration only to the hosts that correspond to the specified role names.
 - Any other hosts that previously had deployed client configurations, but do not have gateway roles assigned to them, will have those client configurations removed from them. This is the new behavior.
- The behavior of the cluster level `deployClientConfig` command, and calling the service level command with no arguments, is unchanged. The command still deploys a new client configuration to all hosts with roles corresponding to the specified service or cluster.
- As this change is due to internal functional changes inside CM, it is not restricted to any new API level. The `deployClientConfig` command in all API levels is affected.

- **Configuration**

- **NameNode configuration** - The decommissioning parameters `dfs.namenode.replication.max-streams` and `dfs.namenode.replication.max-streams-hard-limit` are now available.
- **Hue debug options** - Two service-level configuration parameters have been added to the Hue service to enable Django debug mode and debugging of internal server error responses.

What's New in Cloudera Manager 5.2.7

An issue has been fixed. See [Issues Fixed in Cloudera Manager 5.2.7](#) on page 2667.



Note: Although there is a Cloudera Manager 5.2.7 release, there is no synchronous CDH 5.2.7 release.

What's New in Cloudera Manager 5.2.6

A number of issues have been fixed, see [Issues Fixed in Cloudera Manager 5.2.6](#) on page 2668.

What's New in Cloudera Manager 5.2.5

A number of issues have been fixed, see [Issues Fixed in Cloudera Manager 5.2.5](#) on page 2669.

What's New in Cloudera Manager 5.2.4

There are no changes for Cloudera Manager 5.2.4. It was released to provide the Cloudera Navigator fix in .



Note: Although there is a CDH 5.2.3 release, there is no synchronous Cloudera Manager 5.2.3 release.

What's New in Cloudera Manager 5.2.2

- **HDFS Decommissioning** - The following decommissioning properties have been exposed in Cloudera Manager 5.2.2.
 - **Maximum number of replication threads on a Datanode** (`dfs.namenode.replication.max-streams`)
 - **Hard limit on the number of replication threads on a Datanode** (`dfs.namenode.replication.max-streams-hard-limit`)
- New icons for the KMS and Key Trustee services.

What's New in Cloudera Manager 5.2.1

This release fixes the “POODLE” vulnerability and a number of other issues. See [Issues Fixed in Cloudera Manager 5.2.1](#) on page 2669.

- The YARN `yarn.nodemanager.recovery.dir` property can be configured.
- A health check indicates whether the HDFS metadata upgrade has not been finalized.

What's New in Cloudera Manager 5.2.0

- **OS and database support** - Adds support for Ubuntu Trusty (version 14.04) and PostgreSQL 9.3.
- **Services** - the following new services have been added:
 - **Isilon** - supports the EMC Isilon distributed filesystem.
 - **KMS** - the Java keystore-based key management server.
 - **Key Trustee** - the enterprise-grade key management server using Cloudera Navigator Key Trustee.
 - **Spark** - running Spark applications on YARN. The existing Spark service has been renamed Spark (Standalone).
- **Accumulo** - Kerberos authentication is now supported. If you have been using advanced configuration snippets (safety valves) to configure Kerberos with Accumulo, you may now remove those settings and have Cloudera Manager generate the principal and keytab file for you.
- **HDFS Data at Rest Encryption** -



Note: Cloudera provides the following two solutions for data at rest encryption:

- **Navigator Encrypt** - Is production ready and available for Cloudera customers licensed for Cloudera Navigator. Navigator Encrypt operates at the Linux volume level, so it can encrypt cluster data inside and outside HDFS. Talk to your Cloudera account team for more information about this capability.
- **HDFS Encryption** - Included in CDH 5.2.0 and operates at the HDFS folder level, enabling encryption to be applied only to HDFS folders where needed. This feature has several known limitations. Therefore, Cloudera does not currently support this feature in CDH 5.2 and it is *not* recommended for production use. To try the feature, upgrade to the latest version of CDH 5.

HDFS now implements transparent, end-to-end encryption of data read from and written to HDFS by creating encryption zones. An encryption zone is a directory in HDFS with all of its contents, that is, every file and subdirectory in it, encrypted. You can use either the **KMS** or the **Key Trustee** service to store, manage, and access encryption zone keys.

- **HBase** - Support for configuring hedged reads has been added for HBase. The default configuration is to turn hedged reads off. Cloudera Manager will emit two properties, `dfs.client.hedged.read.threadpool.size` (default: 0) and `dfs.client.hedged.read.threshold.millis` (default: 500ms) to `hbase-site.xml`.

- **ZooKeeper** - the RMI port can be configured. The port is configured using the JDK7 flag `-Dcom.sun.management.jmxremote.rmi.port`. The default value is set to be same as the JMX Agent port. Also, a special value of 0 or -1 disables the setting and a random port is used. The configuration has no effect on versions lower than Oracle JDK 7u4.
- **Cloudera Manager Agent configuration**
 - The supervisord port can now be configured in the Agent configuration `supervisord_port`. The change takes effect the next time supervisord is restarted (not simply when the Agent is restarted).
 - Added an Agent configuration `local_filesystem_whitelist` that allows configuring the list of local filesystems that should always be monitored.
- **Proxy user configuration**
 - All services' proxy user configuration properties have been moved to the HDFS service. Other services running on the cluster inherit the configuration values provided in HDFS. If you have previously configured a service to have values different from those configured in HDFS, then the proxy user configuration properties will be moved to that service's Advanced Configuration Snippet (Safety Valve) for `core-site.xml` to retain existing behavior.

Oozie and Solr are exceptions to this. Oozie proxy user configuration properties have been moved to **Oozie Server Advanced Configuration Snippet (Safety Valve)** for `oozie-site.xml` if they differ from HDFS. Solr proxy user configuration properties have been moved to **Solr Service Environment Advanced Configuration Snippet (Safety Valve)** if they differ from HDFS.
- **Resource management** - YARN and Llama integrated resource management and Llama high availability wizard.
- **New and changed user roles** - BDR Administrator, Cluster Administrator, Navigator Administrator, and User Administrator. The Administrator role has been renamed Full Administrator.
- **Configuration UI**
 - Cluster-wide configuration - you can view all modified settings and configure log directories, disk space thresholds, and port settings.
 - New configuration layout - the new layout provides an alternate way to view configuration pages. In the **classic** layout, pages are organized by role group and categories within the role groups. The **new** layout allows you to filter on configuration status, category, and scope. On each configuration page you can easily switch between the classic and new layout.



Important: The classic layout is the default. All the configuration procedures described in the Cloudera Manager documentation assume the classic layout.

What's New in Cloudera Manager 5.1.6

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.1.6](#) on page 2672.

What's New in Cloudera Manager 5.1.5

A number of issue have been fixed. See [Fixed Issues in Cloudera Manager 5.1.5](#) on page 2673.

What's New in Cloudera Manager 5.1.4

A number of issues have been fixed. See [Fixed Issues in Cloudera Manager 5.1.4](#) on page 2673.

What's New in Cloudera Manager 5.1.3

A number of issues have been fixed. See [Fixed Issues in Cloudera Manager 5.1.3](#).

- **JDK Installation**
 - Users who are adding or upgrading hosts can now choose not to install the JDK that ships with Cloudera Manager.

Cloudera Manager 5 Release Notes

What's New in Cloudera Manager 5.1.2

A number of issues have been fixed. See [Fixed Issues in Cloudera Manager 5.1.2](#).

- **New SAML configuration option**

- You can now specify the binding protocol to be used for AuthNResponses sent from the IDP to Cloudera Manager. Previously, Cloudera Manager would only use HTTP-Artifact, but it is now possible to choose HTTP-Post. HTTP-Artifact remains the default binding.

What's New in Cloudera Manager 5.1.1

An issue has been fixed. See [Issues Fixed in Cloudera Manager 5.1.1](#) on page 2675.

What's New in Cloudera Manager 5.1.0



Important: Cloudera Manager 5.1.0 is no longer available for download from the Cloudera website or from archive.cloudera.com due to the JCE policy file issue described in the [Fixed Issues in Cloudera 5.1.1](#) section of the Release Notes. The download URL at archive.cloudera.com for Cloudera Manager 5.1.0 now forwards to Cloudera Manager 5.1.1 for the RPM-based distributions for Linux RHEL and SLES.

- **SSL Encryption**

- Supports several new SSL-related configuration parameters for HDFS, MapReduce, YARN and HBase, which allow you to configure and enable encrypted shuffle and encrypted web UIs for these services.
- Cloudera Manager now also supports the monitoring of HDFS, MapReduce, YARN, and HBase when SSL is enabled for these services. New configuration parameters allow you to specify the location and password of the truststore used to verify certificates in HTTPS communication with CDH services and the Cloudera Manager Server.

- **Sentry Service**

- A new Sentry service that stores the authorization metadata in an underlying relational database and allows you to use Grant/Revoke statements to modify privileges.
- You can also configure the Sentry service to allow Pig, MapReduce, and WebHCat queries access to Sentry-secured data stored in Hive.

- **Kerberos Authentication**

- Now supports a Kerberos cluster using an Active Directory KDC.
- New wizard to enable Kerberos on an existing cluster. The wizard works with both MIT KDC and Active Directory KDC.
- Ability to configure and deploy Kerberos client configuration (`krb5.conf`) on a cluster.

- **Spark Service** - added the History Server role

- **Impala** - added support for Llama ApplicationMaster High Availability

- **User Roles** - there are two new roles: Operator and Configurator that support fine-grained access to Cloudera Manager features.

- **Monitoring**

- Updates to Oozie monitoring
- New Hive metastore canary

- **UI** - The UI has been updated to improve scalability. The **Home > Status** tab can be configured to display clusters in a full or summary format. There is a new Cluster page for each cluster. The Hosts and Instances pages have added faceted filters.

What's New In Cloudera Manager 5.0.7

A number of issues have been fixed. See [Fixed Issues in Cloudera Manager 5.0.7](#) on page 2678.

What's New in Cloudera Manager 5.0.6

A number of issues have been fixed. See [Fixed Issues in Cloudera Manager 5.0.6](#) on page 2679.

What's New in Cloudera Manager 5.0.5

A number of issues have been fixed. See [Fixed Issues in Cloudera Manager 5.0.5](#) on page 2679.

What's New in Cloudera Manager 5.0.2

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.0.2](#) on page 2679.

What's New in Cloudera Manager 5.0.1

A number of issues have been fixed. See [Issues Fixed in Cloudera Manager 5.0.1](#) on page 2679.

- **Monitoring**

- The Java Garbage Collection Duration health test for the Service Monitor, Host Monitor, and Activity Monitor has been replaced with the new Java Pause Duration health test.

What's New in Cloudera Manager 5.0.0

- **Service and Configuration Management**

- **HDFS** - cache management

- **Resource Management** - Impala admission control

- **Monitoring**

- Host disks overview
- Impala best practices
- HBase table statistics
- HDFS cache statistics

What's New in Cloudera Manager 5.0.0 Beta 2

- **Service and Configuration Management**

- **HDFS**

- HDFS NFS Gateway role
- Supports restoration of HDFS data from a snapshot

- **YARN**

- YARN Resource Manager High Availability
- Resource pool scheduler

- Support for Spark service
- Support for Accumulo service
- Support for service extensibility
- Support to set up Oozie server High Availability
- Granular configuration staleness UI
- Support for setting maximum file descriptors

- **Monitoring**

- Support for monitoring the Cloudera Search/Solr service
- New "failed" and "killed" badges displayed for unsuccessful YARN applications
- More attributes available for filtering displays of YARN applications and Impala queries
- New operational reports added for HBase tables and namespaces, Impala queries, and YARN applications
- Support for creating user-defined triggers for metrics accessible via charts/tsquery



Important: Because triggers are a new and evolving feature, backward compatibility between releases is not guaranteed at this time.

- Charting improvements
 - New table chart type
 - New options for displaying data and metadata from charts
 - Support for exporting data from charts to CSV or JSON files
- **Administrative Settings**
 - Added a new role type with limited administrator capabilities.
 - Cloudera Manager Server and all JVMs will create a heap dump if they run out of memory.
 - Configure the location of the parcel directory and specify whether and when to remove old parcels from cluster hosts.

What's New in Cloudera Manager 5.0.0 Beta 1

- **CDH Version**
 - Supports both CDH 4 and CDH 5
 - CDH 4 to CDH 5 upgrade wizard
 - Support for YARN as a production execution environment
 - MapReduce (MRv1) to YARN (MRv2) configuration import
 - YARN-based resource management for Impala 1.2
- **JDK Version** - Cloudera Manager 5 supports and installs both JDK 6 and JDK 7.
- **Resource Management**
 - Static and dynamic partitioning of resources: provides a wizard for configuring static partitioning of resources (cgroups) across core services (HBase, HDFS, MapReduce, Solr, YARN) and dynamic allocation of resources for YARN and Impala.
 - Pool, resource group, and queue administration for YARN and Impala.
 - Usage monitoring and trending.
- **Monitoring**
 - YARN service monitoring
 - YARN (MRv2) job monitoring
 - Configurable histograms of Impala query and YARN job attributes that can be used to quickly filter query and application lists
 - Scalable back-end database for monitoring metrics
 - Charting improvements
 - New chart types: histogram and heatmap
 - New scale types: logarithmic and power
 - Updates to tsquery language: new attribute values to support YARN and new functions to support new chart types
- **Extensibility**
 - Ability to manage both ISV applications and non-CDH services (for example, Accumulo, Spark, and so on)
 - Working with select ISVs as part of Beta 1
- **Single Sign-On** - Support for SAML to enable single sign-on
- **Parcels**

- Dependency enforcement to ensure incompatible parcels are not used together
- Option to not cache downloaded parcels, to save disk space
- Improved error reporting for management operations
- **Backup and Disaster Recovery (BDR)**
 - HBase and HDFS snapshots: Supports scheduling snapshots on a recurring basis.
 - Support for YARN (MRv2): Replication jobs can now run using YARN (MRv2) instead of MRv1.
 - Global replication page: All scheduled snapshots (HDFS and HBase) and replication jobs for either HDFS or Hive are shown on a single Replications page.
- **Other**
 - Global Search box
 - Several usability improvements
 - Comprehensive detection of configuration changes that require service restarts, refresh and redeployment of client configurations

Incompatible Changes in Cloudera Manager 5

The following sections describe incompatible changes in each Cloudera Manager 5 release.

Incompatible Changes Introduced in Cloudera Manager 5.15.0

- **Oozie Load Balancer Configuration Parameter -**

The `oozie_load_balancer` Cloudera Manager configuration parameter has been changed. Previously, it was specified as `<hostname>:<port>` format. In Cloudera Manager 5.15 and later, the format is `<hostname>`. This format change is incompatible.

Any client reading this value through the API should also read as necessary the load balancer port configuration parameter: `oozie_load_balancer_http_port` or `oozie_load_balancer_https_port`. The correct port parameter depends on whether TLS/SSL is enabled. Check the `oozie_use_ssl` property if you are unsure.

Incompatible Changes Introduced in Cloudera Manager 5.5.0

- Cloudera Manager no longer supports JDK 1.6.

Incompatible Changes Introduced in Cloudera Manager 5.4.0

- The Blacklisted Products property has been removed from the Hosts > Parcels configuration.

Incompatible Changes Introduced in Cloudera Manager 5.3.0

- Oozie metrics - The Oozie metrics framework is now controlled by the **Enable The Metrics Instrumentation Service** flag, which is enabled by default. When enabled, the old 'instrumentation' REST end-point is disabled and metrics are available on the new 'metrics' REST end-point (`hostname:port/v2/admin/metrics`).

Incompatible Changes Introduced in Cloudera Manager 5.2.0

- Due to various internal changes to configuration generation, all service and client configurations will be stale after upgrade. To propagate the updates, restart the cluster and redeploy client configurations.

Incompatible Changes Introduced in Cloudera Manager 5.1.0

- The Limited Administrator role has been renamed Limited Operator. The Limited Operator role is no longer available in Cloudera Manager Express. If you upgrade a Cloudera Manager Express installation, users in the Limited Operator role will not be able to log in. A user in the Administrator role must assign the Read-Only or Administrator role to those users.

Incompatible Changes Introduced in Cloudera Manager 5.0.0

- **Cloudera Manager API**

- New [upgradeCdh](#) command, which upgrades CDH cluster versions. Use this command to upgrade clusters from CDH 4 to CDH 5. The `upgradeServices` command previously used to upgrade CDH cluster versions is no longer supported.
- The `hostId` field now contains a unique UUID and no longer matches the `hostName` field. When referring to a host, both `hostId` and `hostName` are accepted. However, any API clients that were previously cross-referencing host records with external information by `hostName`, but were using the `hostId` field in the API, must be updated to use the `hostName` field. Clients updated in this manner will function correctly with older versions of Cloudera Manager because the `hostName` field has always been present.
- The `clusterName` field displayed when viewing service and role references is now an internal name and may not match the external `displayName` field of the cluster.

- All CDH 5 versions of Hue work only with the default system Python version of the operating system it is being installed on. For example, on RHEL/CentOS 6, you need Python 2.6 to start Hue.
- Cloudera Manager 5.0 includes a change to the value of the `snmpTrapOID`. Earlier releases set the value of `snmpTrapOID` (OID: `.1.3.6.1.6.3.1.1.4.1.0`) wrongly to `clouderaManagerMIBNotifications` (OID `.1.3.6.1.4.1.38374.1.1.1`). This is fixed in Cloudera Manager 5.0 with the correct value, which is `clouderaManagerAlert` (OID `.1.3.6.1.4.1.38374.1.1.1.1`). This change will break SNMP server setups that are configured to expect `clouderaManagerMIBNotifications`. Cloudera Manager administrators should configure their SNMP receivers to accept the corrected OID.
- The default values for the following configurations have changed to include the JVM option `-Djava.net.preferIPv4Stack=true`, which sets the preferred protocol stack to IPv4 on dual-stack machines. Any values set to the old defaults will automatically be changed to the new default when upgrading to Cloudera Manager 5.

- MapReduce client configuration:

- `hadoop-env.sh`: added to `HADOOP_CLIENT_OPTS`
- `mapred-site.xml`: added to `mapred.child.java.opts`

- YARN client configuration:

- `hadoop-env.sh`: added to `YARN_OPTS`
- `mapred-site.xml`: added to `yarn.app.mapreduce.am.command-opts`, `mapreduce.map.java.opts`, and `mapreduce.reduce.java.opts`

- HDFS client configuration: `hadoop-env.sh`: added to `HADOOP_CLIENT_OPTS`

- Hive client configuration: `hive-env.sh`: added to `HADOOP_CLIENT_OPTS`

- MapReduce health tests have been removed:

- Job failure
- Map backlog
- Reduce backlog
- Map locality

If needed, the test can be replaced with a trigger. For example:

- Looks at all the jobs that completed in the last hour and if there are more than 10% of failed jobs, change the health of the service to concerning:

```
IF (select (jobs_failed_rate * 3600) as jobs_failed,
((jobs_failed_rate + jobs_completed_rate + jobs_killed_rate) * 3600)
as all_jobs where roleType=JOBTRACKER AND serviceName=$SERVICENAME
and last(jobs_failed_rate / (jobs_failed_rate + jobs_completed_rate +
jobs_killed_rate)) >= 10 ending at $END_TIME duration "PT3600S")
DO health:concerning
```


- If there are more than 50% maps waiting than total slots available, health goes concerning.

```
IF (select waiting_maps / map_slots where roleType=JOBTRACKER and serviceName=$SERVICENAME
and last(waiting_maps / map_slots) > 50)
DO health:concerning
```

- If there are more than 50% reduce waiting than total slots available, health goes concerning.

```
IF (select waiting_reduces / reduce_slots where roleType=JOBTRACKER and
serviceName=$SERVICENAME
and last(waiting_reduces / reduce_slots) > 50)
DO health:concerning
```

- HDFS checkpointing metrics have been removed:

- end_checkpoint_num_ops
- end_checkpoint_avg_time
- start_checkpoint_num_ops
- start_checkpoint_avg_time

Incompatible Changes Introduced in Cloudera Manager 5.0.0 Beta 2

- Impala releases earlier than 1.2.1 are no longer supported.
- Some of the constants identifying health tests have changed. The following existed in Cloudera Manager 4:

- FAILOVERCONTROLLER_FILE_DESCRIPTOR
- FAILOVERCONTROLLER_HOST_HEALTH
- FAILOVERCONTROLLER_LOG_DIRECTORY_FREE_SPACE
- FAILOVERCONTROLLER_SCM_HEALTH
- FAILOVERCONTROLLER_UNEXPECTED_EXITS

They are now:

- MAPREDUCE_FAILOVERCONTROLLER_FILE_DESCRIPTOR
- MAPREDUCE_FAILOVERCONTROLLER_HOST_HEALTH
- MAPREDUCE_FAILOVERCONTROLLER_LOG_DIRECTORY_FREE_SPACE
- MAPREDUCE_FAILOVERCONTROLLER_SCM_HEALTH
- MAPREDUCE_FAILOVERCONTROLLER_UNEXPECTED_EXITS

and

- HDFS_FAILOVERCONTROLLER_FILE_DESCRIPTOR
- HDFS_FAILOVERCONTROLLER_HOST_HEALTH
- HDFS_FAILOVERCONTROLLER_LOG_DIRECTORY_FREE_SPACE
- HDFS_FAILOVERCONTROLLER_SCM_HEALTH
- HDFS_FAILOVERCONTROLLER_UNEXPECTED_EXITS

The reason for the change is to better distinguish between MapReduce and HDFS failover controller monitoring in the health system.

Incompatible Changes Introduced in Cloudera Manager 5.0.0 Beta 1

- **Services**

- **Impala** - With Cloudera Manager 4.8 (released in late November 2013), only Impala 1.2.1 is supported, due to the introduction of the Impala Catalog Server. However, CDH 5.0.0 Beta 1 was released with Impala 1.2.0 (Beta). Therefore, if you upgrade from Cloudera Manager 4.8 (with Impala 1.2.1) to Cloudera Manager 5.0.0 Beta 1, and then upgrade your CDH to CDH 5.0.0 Beta 1, your version of Impala will be downgraded to Impala 1.2.0 from 1.2.1. This will result in some loss of functionality. See [New Features in Impala](#) for a list of the new features in Impala 1.2.1 that are not in Impala 1.2.0 (Beta).

- **Hive** - HiveServer 2 is a mandatory role for Hive in CDH 5.
- **Hue** - In CDH 5, Hue no longer has a Beeswax Server role. Hue now submits queries to HiveServer2.
- **HDFS** - Cloudera Manager 5 does not support NFS-mounted shared edits directories for HDFS High Availability. It only supports the Quorum Journal method for shared edits. If you upgrade from Cloudera Manager 4 with a working CDH 4 High Availability configuration that uses NFS-mounted directories, your installation will continue to work until you disable High Availability. You will not be able to re-enable High Availability with NFS-mounted directories. Furthermore, you will not be able to upgrade to CDH 5 unless you disable High Availability, and you will need to use Quorum-based storage in order to re-enable High Availability after the upgrade.
- **YARN**
 - The YARN (MRv2) configuration `mapreduce.job.userlog.retain.hours` has been replaced by `yarn.log-aggregation.retain-seconds`. Any existing value in `mapreduce.job.userlog.retain.hours` will be lost. However, this configuration never had any effect, so no functionality is affected.
 - The following configuration parameters were removed from YARN. These never had any effect, so no functionality is affected.
 - `mapreduce.jobtracker.maxtasks.perjob`
 - `mapreduce.jobtracker.handler.count` (non-functional duplicate of `yarn.resourcemanager.resource-tracker.client.thread-count`)
 - `mapreduce.jobtracker.persist.jobstatus.active`
 - `mapreduce.jobtracker.persist.jobstatus.hours`
 - `mapreduce.job.jvm.numtasks`
 - The following YARN configuration parameters were replaced. Only the YARN parameters were replaced. Old configurations will be lost, but they never had any effect so this does not affect functionality.
 - `mapreduce.jobtracker.restart.recover` replaced by `yarn.resourcemanager.recovery.enabled` (changed from Gateway to ResourceManager)
 - `mapreduce.tasktracker.http.threads` replaced by `mapreduce.shuffle.max.connections`
 - `mapreduce.jobtracker.staging.root.dir` replaced by `yarn.app.mapreduce.am.staging-dir`
 - Cloudera Manager 5 sets the default YARN Resource Scheduler to FairScheduler. If a cluster was previously running YARN with the FIFO scheduler, it will be changed to FairScheduler the next time YARN restarts. The FairScheduler is only supported with CDH 4.2.1 and later, and older clusters may hit failures and need to manually change the scheduler to FIFO or CapacityScheduler. See the Known Issues section of this Release Note for information on how to change the scheduler back to FIFO or CapacityScheduler.

Changed Features and Behaviors in Cloudera Manager 5

The following sections describe what's changed in each Cloudera Manager 5 release.



Note: Rolling upgrade is not supported between CDH 4 and CDH 5. Rolling upgrade will also *not* be supported from CDH 5.0.0 Beta 2 to any later releases, and may not be supported between any future beta versions of CDH 5 and the General Availability release of CDH 5.

What's Changed in Cloudera Manager 5.16

- **Backup and Disaster Recovery** - BDR now ignores Kudu tables during replication. The change does not affect functionality since BDR does not support Kudu tables. This change was made to guard against data loss due to how the Hive MetaStore, Imapla, and Kudu interact.

What's Changed in Cloudera Manager 5.15

- **Host Inspector** - Previously, the Host Inspector displayed a list of components and version numbers. Now, all mismatched versions of a component are grouped together.
- **Key Trustee Server** - You cannot run the CDH Upgrade wizard from Key Trustee Server clusters. Previously, you could run the upgrade wizard on a Key Trustee Server cluster even though it is not part of the CDH upgrade.
To upgrade Key Trustee Server, distribute and activate a new Key Trustee Server parcel. Then restart the Key Trustee Server service. OPSAPS-42792
- **Key Trustee Server KMS Service** - Previously, you always had to enter a key admin user when you add a Key Trustee Service. Now, as long as the key admin user or key admin group contain a non-empty value, the Generate ACLs button is enabled.
- **Menu and UI Labels**
Previously, **Hive-on-S3** was a **Replication Option** when you create or edit a Hive replication schedule. The option has been renamed **Hive-on-Cloud** since Microsoft ADLS is now supported by Backup and Disaster Recovery
- **Stale Services** - Previously, stale services were not preselected in the Restart Staled Services wizard. They are now preselected.

What's Changed in Cloudera Manager 5.14

- **API**
 - The enum field `apiCluster.apiClusterVersion` is deprecated. Use the string field `apiCluster.fullVersion` instead.
 - You can now configure the timeout for `TimeSeriesQueryService` with the API. Previously, there was a hard-coded 20 second timeout.
- **Authorization**
 - The Reports Manager role is no longer required to run Cloudera Management Service anymore.
 - When Cloudera Manager users external authentication, such as LDAP, role assignment from the **Users** page is now disabled.
- **BDR**
 - You no longer have to select the **Skip checksum check** option for BDR replication to function properly when one or both clusters use encryption. The checksum check is transparently and automatically skipped in this situation.
 - Removed an extra RPC call to source NameNode in the mapper phase of back up replications. The fix reduces the number of RPC calls to the NameNode to less than half. After this fix:
 - Customers will see an improvement in replication (Backup and Disaster Recovery) performance.
 - Customers will be better able to scale their replication jobs and the production cluster NameNode is no longer a bottleneck during the mapper phase.
- **Installation**
 - CDH4 no longer appears as an install option in the Cloudera Manager Installation Wizard. CDH4 builds will also not be recognized if users try to use the REST API to install CDH4 on a new host.
 - The list of available versions of CDH5 for package-based installation is now obtained by querying `archive.cloudera.com`. New maintenance versions of CDH will now be available for install without having to upgrade Cloudera Manager. If `archive.cloudera.com` is not reachable, the list of CDH versions available will be empty.
 - Cloudera Manager now only installs Java 7. Previously when you installed Java on managed hosts using Cloudera Manager, both Java 6 and 7 were installed.

What's Changed in Cloudera Manager 5.13

- **External Accounts**

The **AWS Credentials** and **Altus Credentials** under **Administration** have been consolidated. You can find the pages for the credentials by selecting **Administration > External Accounts**.

- **RPC Wait Behavior**

When starting or restarting HDFS, Cloudera Manager now waits for any NameNode daemons that are started to begin responding to RPC requests before considering the start/restart operation to be complete. Previously, this wait was only performed in certain workflows; now, it is always done when starting HDFS.

- **Hue Load Balancer Enabled by Default**

For better Hue web site performance, the Hue Load Balancer is now enabled by default. In a secure cluster, Hue Load Balancer (apache httpd 2.4) requires "use_x_forwarded_host" to be set to "true". This change will cause staleness.

What's Changed in Cloudera Manager 5.12

- **Backup and Disaster Recovery**

- **Replicate Impala metadata feature for CDH 5.12 or later**

Replication of "native" Impala UDFs (transient UDFs created with the old syntax prior to CDH 5.7) is no longer available for CDH clusters that run version 5.12 or later because Impala native UDFs are deprecated. If you still use Impala native UDFs and want them to be replicated, you should recreate them within the Hive Metastore using the new Create Function syntax supported since CDH 5.7.

- **Dynamic Resource Pool Scheduling Rules**

You can no longer specify a one-time scheduling rule for a dynamic resource pool. Recurring scheduling rules are still supported. To make one-time changed to resource pool configuration, simply update the pools via the Cloudera Manager UI or API.

- **Add Cluster Wizard**

MapReduce1 has been removed as an option from the **Add Cluster Wizard**. It can still be added to a cluster after initial cluster creation or through the API.

- **Web UI**

- **Icons**

The icons in the Cloudera Manager web UI are more usable and distinct.

- **Stack Traces**

By default, stack traces are no longer displayed in the web UI. Please view the Cloudera Manager Server log under **Diagnostics > Server Log** to view stack traces of exceptions generated by Cloudera Manager.

- **Key Trustee Parcel**

The Key Trustee Parcel will no longer be released via archive.cloudera.com. The parcel will now be released via <http://www.cloudera.com/downloads>. Parcels already released on the archive site will continue to be available there.

- **Database names and DB usernames for the Cloudera Manager database.**

You can now only use alphanumeric characters and underscores for database names and usernames with the `scm_prepare_database.sh` script to ensure properly supported database names and usernames.

- **Cloudera Manager Agent Dependencies**

The Cloudera Manager Agent has new package dependencies. The `netstat` and `ifconfig` commands have been replaced with the `ss` and `ip` commands respectively.

For more information, see [iproute package requirement](#).

What's Changed in Cloudera Manager 5.11

- **Change to default fencing method for HDFS High Availability**

The default value of the HDFS High Availability Fencing Methods property (`dfs.ha.fencing.methods`) has been changed from `shell(/cloudera_manager_agent_fencer.py)` to `shell(true)`. While the old default value was reasonable for early versions of HDFS, it is no longer recommended for HAQJ-based versions of HDFS HA, as it can lead to HDFS service outages. The new default, `shell(true)`, is the setting that Cloudera recommends. It uses the built-in HDFS fencing method, which causes the fenced NameNode to exit if it attempts a write operation when it is not supposed to be active. If auto-restart is enabled for the NameNode, Cloudera Manager will then restart it.

Most clusters already have the HDFS High Availability Fencing Methods explicitly set to `shell(true)` since that is the value set by the HDFS Enable High Availability Wizard. When such clusters are upgraded to Cloudera Manager 5.11, this explicit setting of `shell(true)` will be removed. This does not change the effective value of the property since it takes on the new default value, which is also `shell(true)`. The cluster will experience no change in fencing behavior.

When clusters that have HDFS High Availability Fencing Methods set to the pre-5.11 default value are upgraded to Cloudera Manager 5.11, the effective value of the property changes to `shell(true)`, and the fencing behavior changes accordingly. This also causes the HDFS service to become stale, requiring a restart of HDFS and dependent services. If HDFS High Availability Fencing Methods is set to a non-default value other than `shell(true)`, no change occurs in the property's value when the cluster is upgraded to Cloudera Manager 5.11. Because the fencing method `shell(/cloudera_manager_agent_fencer.py)` can lead to service outages, a new configuration warning message will be displayed if it is in use after the upgrade to Cloudera Manager 5.11 is complete.

What's Changed in Cloudera Manager 5.10

- **Configuration Changes**

- **Llama configuration options removed**

Llama roles must be removed before upgrading to CDH 5.10.0 or higher. If your cluster has an Impala Llama role and you are using Cloudera Manager to upgrade, Cloudera Manager displays an error message and prevents the upgrade from going forward. You must first remove any existing Llama roles, using the **Disable YARN and Impala Integrated Resource Management** command (Before upgrading, go to the Impala Service and select **Actions > Disable YARN and Impala Integrated Resource Management**.) Cloudera Manager also generates a configuration error if a Llama role is added to a CDH 5.10 cluster, and prevents Impala from starting until the role is removed.

- **JAVA_HOME in agent configuration deprecated in 5.x**

Specifying host agent environment variable `CMF_AGENT_JAVA_HOME` is deprecated and will not be supported in a future release. Instead, specify **Java Home Directory** property in Cloudera Manager.

- **The `hbase.client.scanner.timeout.period` property is now configurable for clients**

The Scan API-related property `hbase.client.scanner.timeout.period` in the HBase service is now configurable at the Client (Gateway) level. Its value may be set to be equal to or lesser than the RegionServer-side equivalent of **HBase RegionServer Lease Period**.

- **Warning if KeyTrustee Server is located on cluster machines**

The KeyTrustee Server should be installed on dedicated hosts in a separate cluster. While this is adequate for proof of concept deployments, Cloudera does not recommend such installations for use in a production environment. Cloudera now shows a warning to the user when adding a Key Trustee server service onto an existing CDH cluster host.

What's Changed in Cloudera Manager 5.9.0

- You must now have the **BDR, Cluster, or Full Administrator** role to view the following pages:
 - HDFS File Browser
 - Directory Usage Report
 - HBase Table Browser
 - Solr Collections
 - HBase Table Statistics

Previously, you only needed the **Read-Only** role.

What's Changed in Cloudera Manager 5.8.0

- The YARN service's list of **Allowed System Users** now includes the `hbase` user by default. The reason for this change is that several essential HBase tools such as the MOB Sweeper, Import/Export tools, and CopyTable, need to interact with HBase as the `hbase` user to be able to execute MapReduce jobs.

Note that this change is only applicable to new Cloudera Manager deployments. Upgrading to Cloudera Manager 5.8 will not add the `hbase` user to the list of defaults.


What's Changed in Cloudera Manager 5.7.0

- The Navigator Metadata Server requires 192 MiB of Java PermGen space instead of 128 MiB. The value of this internal setting used by the JDK is increased automatically when upgrading to Cloudera Manager 5.7.
- The default value for `hive.compute.query.using.stats` is changed to `false`. The reason for the change is that certain queries such as `count`, `max`, and `min` return incorrect results with this optimization on.
- By default, Hive sessions now only consider sessions with no recent activity to be idle (`hive.server2.idle.session.timeout_check_operation`) and idle session timeouts have been reduced (`hive.server2.idle.session.timeout` and `hive.server2.idle.operation.timeout`). This helps reduce the strain on HiveServer2 from too many open sessions.
- Cloudera Manager no longer automatically refreshes scheduler configurations when dynamic resource pool settings are changed. You must explicitly refresh the configurations. This allows you to schedule the changes to minimize the impact on your cluster.
- For YARN, the default number of log directories (`yarn.nodemanager.log-dirs`) has changed from 1 to be equal to the number of mount points, to prevent applications with a large number of logs from filling up a single disk.
- The default for Java Heap Size of JournalNode in Bytes is now 512 MB.
- The **Sources** page for HDFS and Hive replications has been removed. A list of sources is available from a drop-down menu when you schedule a replication.
- The number of watched directories you can specify for the Disk Usage Report is now unlimited.
- Cloudera Manager now uses a new memory allocation algorithm to allocate memory when multiple roles are installed on the same host.
- User sessions in the Cloudera Manager Admin Console now timeout after a configurable period of time of inactivity. A dialog box warns the user before automatically logging out the user.
- The **All Recent Commands** page now loads more quickly.
- The **Disk Usage** reports now have links that take users to the **Directory Usage Report** with the correct filter applied.
- When searching for hosts on the **Hosts** page, you can now filter the hosts list by entering search terms (hostname, IP address, or role) in the search box separated by commas or spaces. You can use quotes for exact matches (for example, strings that contain spaces, such as a role name) and brackets to search for ranges. Hosts that match any of the search terms are displayed.
- Isilon is now supported as a source or destination service for HDFS replications.
- For CDH 5.7 and higher if `CDH_PYTHON` is set by a Spark plug-in, `PYSPARK_PYTHON` is set to `CDH_PYTHON` in `spark-env.sh`. If you install a Python runtime parcel, such as the [Anaconda](#) parcel, Python Spark jobs run in both YARN client and YARN cluster modes are automatically configured by redeploying the Spark client configuration.

What's Changed in Cloudera Manager 5.5.0

- Removed `-XX:-CMSConcurrentMTEnabled` from the default JVM options. This setting makes the JVM run in single threaded mode. This was needed for Java 1.6_31 and lower but not for Java 1.6_32 or higher. Anybody using Java 1.6_31 or lower should upgrade to the latest recommended version of Java 1.7.

This change causes all roles to be stale after you upgrade to Cloudera Manager 5.5 and they are indicated as requiring restart in the Cloudera Manager Admin Console. However, as with any upgrade, this is a valid, functional Cloudera Manager state, and the cluster only needs to be restarted when you want the new configurations to take effect.

- `HADOOP_USER_CLASSPATH_FIRST` is now adhered to in Hadoop client configurations. After you upgrade Cloudera Manager, services display a client configuration redeployment required icon .
- For RHEL 7, the `force_start`, `fast_*`, `clean_*`, and `hard_*` commands on the `server-scm-*` services no longer work, as custom start, restart, and stop commands are not supported on systemd based distributions. These have been replaced with `*_next_*` operations, which do not trigger an immediate operation, but signal that the next invoked operation will be forced, fast, clean, or hard.
- The Cloudera EULA is now shown when using the Cloudera Manager Admin Console for the first time.
- The Home tab has been removed from the Cloudera Manager Admin Console navigation bar. You can return to the Home page Status tab by clicking the Cloudera Manager logo.
- All the icons have been refreshed to make them cleaner and easier to read.
- In 5.4.0 an externally assigned role was combined with a Cloudera Manager assigned role and the user had the union of the role privileges. As a consequence, an external user could be assigned an administrator role in Cloudera Manager and they would be an administrator regardless of the externally assigned role. Now only the externally assigned roles are respected. No roles can be assigned to an external user in Cloudera Manager and any roles for an external user in the Cloudera Manager are ignored.

As a result of this change, external users with previously-assigned Cloudera Manager roles will have their permissions modified depending on the LDAP group they belong to. To restore permissions for external users, configure the LDAP groups for these users by navigating to **Administration > Settings**, and click **Category > External Authentication** to display the relevant properties.

- Cloudera Manager and CDH components support TLS 1.0, TLS 1.1, and TLS 1.2, but not SSL 3.0. SSL remains part of the TLS/SSL name for historical reasons.
- The label on the **Generate Credentials** button has been changed to **Generate Missing Credentials** to better reflect the fact that it only creates Kerberos principals that are not present yet in Cloudera Manager.
- Cloudera Manager now downloads binaries from <https://archive.cloudera.com> instead of <https://archive.cloudera.com>.
- The embedded `hbc` feature has been removed from HBase monitoring for stability reasons.
- Increased the default heap sizes for Hive roles. On clusters with sufficient memory, newly created Hive roles have these values:
 - HiveServer2 - 4 G heap, 512 M perm gen
 - Hive Metastore - 8 G heap, 512 M perm gen
 - Gateway - 2 G heap, 512 M perm gen
- By default, Oozie now purges eligible completed workflows and coordinator actions for long-running coordinator jobs.
- Oozie actions that omit the `<job-tracker>` and `<name-node>` elements (and the workflow does not define them in the `<global>` section) use the default values for the JobTracker, Resource Manager, and NameNode from Cloudera Manager in CDH 5.5 and higher.
- Increased the defaults for Oozie parameters:
 - `oozie.service.CallableQueueService.callable.concurrency` - 10
 - `oozie.service.CallableQueueService.threads` - 50

- Sqoop 2 is no longer in the default services to be created in any of the options in the installation wizard. You can choose to add it to the Custom Services option in the Installation wizard or can add it with the Add Service wizard after installation.
- For CDH 5.5.0 and higher the default values of the YARN properties `mapreduce.[map|reduce].java.opts.max.heap` and `mapreduce.[map|reduce].memory.mb` have been changed to 0, which tells YARN to automatically select a default. This helps avoid issues where either `heap` or `memory.mb` is updated, but not the other one (`memory.mb` should be ~30% higher than `heap` to allow for JVM overhead).
- The Host DNS Resolution Duration health test was removed. Its functionality is now covered in the Host DNS Resolution health test.
- The default Replication Strategy is now **Dynamic**.

What's Changed in Cloudera Manager 5.4.1

HDFS Read Throughput Impala query monitoring property is misleading

The `hbase_bytes_read_per_second` and `hdfs_bytes_read_per_second` Impala query properties have been renamed to `hbase_scanner_average_bytes_read_per_second` and `hdfs_scanner_average_bytes_read_per_second` to more accurately reflect that these properties return the average throughput of the query's HBase and HDFS scanner threads respectively. The previous names and descriptions gave the impression that these properties were the query's total HBase and HDFS throughput, which was not accurate.

What's Changed in Cloudera Manager 5.4.0

- Cloudera Manager checks the specified version of CDH before an installation and upgrade to ensure that it is compatible with Cloudera Manager before proceeding. Specifically, for Cloudera Manager 5.4 that means no version of CDH newer than 5.4.x is supported (Cloudera Manager must be upgraded before upgrading to such a version of CDH). Cloudera Manager no longer shows these "too-new" versions of CDH. The 'latest' parcel repository URL will be replaced by the 'latest_supported' repository in the parcel configuration.
- The minimum Java heap size for the Activity Monitor, Host Monitor, and Service Monitor has been changed from 50 MB to 256 MB.
- Regenerating Kerberos principals will be denied if any roles that are using those principals are running. Stop those roles and then attempt to regenerate the principals.
- In previous versions of Cloudera Manager, the 'version' attribute in `tsquery` had values that were integers, for example, 4 for CDH4, 5 for CDH5, -1 for Cloudera Manager. Starting in the Cloudera Manager 5.4, the values for the 'version' attribute are in release string format, for example "cdh5.0.0".
- **Hive**
 - `hive.exec.reducers.max` default value changed from 999 to 1099
 - `hive.exec.reducers.bytes.per.reducer` default value changed from 1 GB to 64 MB
 - The default heap size for the Hive CLI is increased to 1 GB.
 - The property `hive.log.explain.output` is known to create instability of Cloudera Manager Agents in some specific circumstances, specially when the hive queries generate extremely large EXPLAIN outputs. Therefore, the property has been hidden from the Cloudera Manager configuration UI. The property can still be configured through the use of advanced configuration snippets.
- **Impala** - The Impala Daemon now supports the Impala Maximum Log Files property which specifies the total number of log files per severity level that should be retained before they are deleted. By default, after upgrading to CDH 5.4 this property is set to 10, which means that Impala Daemons will only retain up to 10 log files for each severity level. Any additional files will be deleted.
- **HBase** - Moved three settings for HBase coprocessors from Main to Advanced category:
 - Service Wide > HBase Coprocessor Abort on Error: move to 'Service Wide > Advanced > HBase Coprocessor Abort on Error'
 - 'Master Default Group > HBase Coprocessor Master Classes': move to 'Master Default Group > Advanced > HBase Coprocessor Master Classes'
 - RegionServer Default Group > HBase Coprocessor Region Classes': move to 'RegionServer Default Group > Advanced > HBase Coprocessor Region Classes'

What's Changed in Cloudera Manager 5.3.2

- Turning on the internal HBase canary (not to be confused with Cloudera Manager monitoring canary) is optional. On new clusters, it will not be enabled by default. Existing clusters will continue to run the canary until it is disabled from the HBase configuration page.

What's Changed in Cloudera Manager 5.3.0

- **Cloudera Manager upgrade** - If you have any active commands running before upgrade, the server *will fail to start* after upgrade. This includes commands a user might have run and also for commands Cloudera Manager automatically triggers, either in response to a state change, or something that's on a schedule.

What's Changed in Cloudera Manager 5.2.1

- The default value of the YARN `yarn.nodemanager.recovery.dir` property has changed from `{hadoop.tmp.dir}/yarn-nm-recovery` to `/var/lib/hadoop-yarn/yarn-nm-recovery`.

What's Changed in Cloudera Manager 5.2.0

- **Rolling upgrade** - As a result of a recent change in the way DataNodes handle block deletions during a rolling upgrade ([HDFS-5907](#)), the Trash directory may grow unexpectedly while the upgrade is in progress. Deleted blocks are kept during upgrade in case you want to roll back. The blocks are cleaned up after you finalize the upgrade.
- **Agent** -
 - The `hard_stop`, `hard_restart`, and `clean_restart` commands now show a warning message about the impact of using these commands instead of performing the actions. To actually perform the actions, you use the `hard_stop_confirmed`, `hard_restart_confirmed`, and `clean_restart_confirmed` commands.
 - The default supervisor port is changed from 9001 to 19001
- YARN application attributes renamed: `slot_millis` to `slots_millis` and `fallow_slot_millis` to `fallow_slots_millis`

What's Changed in Cloudera Manager 5.1.0



- **UI refresh for scalability**
- Revised authorization privilege model in Sentry.

What's Changed in Cloudera Manager 5.0.0

- MapReduce now inherits topology from HDFS NameNode. Topology configuration for MapReduce JobTracker was removed. The configuration was redundant and the two parameters should always have been set to the same value.
- **UI**
 - The Clusters tab no longer has Activities, Other, and Manage Resources sections.

What's Changed in Cloudera Manager 5.0.0 Beta 2

- **Product**
 - Cloudera Backup and Disaster Recovery (BDR) is now included with Cloudera Enterprise.
 - Cloudera Standard has been renamed to Cloudera Express.
- **OS and packaging**
 - The name of the Cloudera Manager embedded database package has changed from `cloudera-manager-server-db` to `cloudera-manager-server-db-2`. For details, read the upgrade and install topics for your OS.
 - Support for Ubuntu 10.04 and Debian 6.0 is deprecated.
- **HDFS** - enabling High Availability automatically enables auto-failover, unlike in Cloudera Manager 4 where enable auto-failover was a separate command.

- **HBase**
 - In CDH 5 there is no HBase canary because HBase is now monitored by a watchdog process. In CDH 4, the HBase canary is still used.
 - The RegionServer default heap size has been increased to 4GB.
- **Monitoring**
 - Chart "Views" and actions related to views have been renamed to "Dashboard".
 - Changes to how attribute filters are displayed in the Impala queries and YARN applications screens
 - The outdated configuration indicator on the Home, service, and role pages has a new graphic  and now has a tooltip that displays whether a cluster refresh or restart is required. There is a new indicator  for changes that require redeploying client configurations. You can click an indicator to go to the new Stale Configurations page to view and resolve the conditions that gave rise to the indicator.
 - To match the naming convention of tsquery metrics, multiword Impala query and YARN application attribute names have changed from camel case to using an underscore separator. For example `queryType` has changed to `query_type`. For backward compatibility, camel case names are still supported.
- **UI**
 - The main navigation bar in Cloudera Manager Admin Console has been reorganized. The Services tab has been replaced by a Clusters tab that contains links to individual services, which were previously under the Services tab, Activities and Reports sections, which were removed from the main bar, and a new Manage Resources section, which contains links to the new resource pools and service pools features. The All Services page has been removed.
 - The "Safety Valve" properties have been renamed "Advanced Configuration Snippet".
 - The screen for specifying assignment of roles to hosts has been redesigned for improved scalability and usability.
- **Misc**
 - The `io.compression.codecs` property has moved from MapReduce to HDFS.

What's Changed in Cloudera Manager 5.0.0 Beta 1

- When CDH 5 is installed, YARN is installed by default, rather than MapReduce, and is the default execution environment. MapReduce is deprecated in CDH 5 but is fully supported for backward compatibility through CDH 5. In CDH 4, MapReduce is still the default.
- The setting for `yarn.scheduler.maximum-allocation-mb` has been increased to a default of 64GB.
- The minimum heap size for the Solr service has been increased to 200MB (from 50MB previously) to enable it to better handle collection creation.

Known Issues and Workarounds in Cloudera Manager 5

The following sections describe the current known issues in Cloudera Manager 5.

Cloudera Manager 6.x issue with the service role Resume

If a selected service role on a node is restarted and fails, and the customer clicks the "Resume" button in Cloudera Manager, the service role on all of the nodes will be restarted concurrently.

Products affected: Cloudera Manager

Releases affected:

- Cloudera Manager 5.5 and later
- Cloudera Manager 6.0 until 6.3.3
- Cloudera Manager 7.1.x

Users affected: Users with admin role in Cloudera Manager can impact end users of the service.

Impact: In production clusters this can result in a cluster-wide service outage; Already observed for the YARN service and the HDFS service in a few clusters.

Severity: High

Action required:

- A workaround exists where instead of performing a restart we recommend performing a stop/start of the services.
- Issue is fixed in CM-6.3.4, CM-7.2.1 and above.

Knowledge Article: For the latest update on this issue see the corresponding Knowledge article: [Cloudera Customer Advisory: Cloudera Manager 6.x issue with service role Resume](#)

ZooKeeper JMX did not support TLS when managed by Cloudera Manager

The ZooKeeper service optionally exposes a JMX port used for reporting and metrics. By default, Cloudera Manager enables this port, but prior to Cloudera Manager 6.1.0, it did not support mutual TLS authentication on this connection. While JMX has a password-based authentication mechanism that Cloudera Manager enables by default, weaknesses have been found in the authentication mechanism, and Oracle now advises JMX connections to enable mutual TLS authentication in addition to password-based authentication. A successful attack may leak data, cause denial of service, or even allow arbitrary code execution on the Java process that exposes a JMX port. Beginning in Cloudera Manager 6.1.0, it is possible to configure mutual TLS authentication on ZooKeeper's JMX port.

Products affected: ZooKeeper

Releases affected: Cloudera Manager 6.1.0 and lower, Cloudera Manager 5.16 and lower

Users affected: All

Date/time of detection: June 7, 2018

Severity (Low/Medium/High): 9.8 High ([CVSS:3.0/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H](#))

Impact: Remote code execution

CVE: CVE-2018-11744

Immediate action required: Upgrade to Cloudera Manager 6.1.0 and enable TLS for the ZooKeeper JMX port by turning on the configuration settings “Enable TLS/SSL for ZooKeeper JMX” and “Enable TLS client authentication for JMX port” on the ZooKeeper service and configuring the appropriate TLS settings. Alternatively, disable the ZooKeeper JMX port via the configuration setting “Enable JMX Agent” on the ZooKeeper service.



Note: Disabling the ZooKeeper JMX port prevents Cloudera Manager from performing health checks on the ZooKeeper service.

Addressed in release/refresh/patch: Cloudera Manager 6.1.0

Yarn Application Masters fail when a concurrent Deploy Client Configuration job is executed

Yarn Application Masters fail because the `container-executor.cfg` file is missing when a YARN application is executed concurrently with a Deploy Client Configurations job. The container-executor hierarchy has been moved to `/var/lib/yarn-ce`, which is not modified by the Deploy Client Configuration command.

Cloudera Bug: OPSAPS-24398

No changes are required for Cloudera Manager/CDH 5.16.0 and higher. Use the following **Workaround** for Cloudera Manager/CDH 5.15.1 and 5.14.4:

1. On the Cloudera Manager server host, add the following line to the `/etc/default/cloudera-scm-server` file:

```
export CMF_FF_YARN_SAFE_CONTAINER_EXECUTOR_DIR=true
```

2. Restart the Cloudera Manager Server:

```
service cloudera-scm-server restart
```

3. Restart the NodeManager(s) so they can pick up the new configuration change:

- a. In the Cloudera Manager Admin console, go to the YARN service.
- b. Click the **Instances** tab.
- c. Select all hosts with the NodeManager **Role Type**.
- d. Click **Actions for Selected > Restart**.

Fixed Versions: Cloudera Manager/CDH 5.16.1 5.16.0 5.15.1, 5.14.4

Affected Versions Cloudera Manager 5.15 and lower

Cloudera Manager Agent Failed to deactivate alternatives for parcels

In the Cloudera Manager host's agent logs, during the agent startup there are many repeated error entries if the listed parcel in the error logs has a "Distributed" status on the Cloudera Manager Parcels page. The confirmed issue is that the alternative configuration for the second (deactivated) parcel isn't present, resulting in the error seen in the agent log file.

For more information, see this [knowledge base article](#).

Workaround: Un-distribute/remove the affected parcel(s) from the host.

Cloudera bug: OPSAPS-52745

Fixed versions: Cloudera Manager 6.x

Affected versions: Cloudera Manager 5.x

The client configuration deployment fails after upgrade

After upgrade to Cloudera Manager 5.x, client configuration deployment fails with the following error:
"/var/lib/alternatives/hadoop-conf empty!"

Workaround: Either update the /var/lib/alternatives/hadoop-conf file from other working hosts, or remove it and restart the CM agent. This re-generates the file. For more information, see this [knowledge base article](#).

Cloudera bug: OPSAPS-38704

Fixed versions: Cloudera Manager 6.x

Affected versions: Cloudera Manager 5.x

Backup and Disaster Recovery replication with Isilon storage can fail

Backup and Disaster Recovery (BDR) jobs with Isilon storage can fail due to Cloudera Manager's host selection policy for BDR jobs or Cloudera Manager not accepting Isilon as a valid storage type for HDFS replication.

Affected Versions: Cloudera Manager 5.8.5, 5.9.3, 5.11.1, 5.12, 5.14.x, 5.15.1

Fixed Versions: Cloudera Manager 5.15.2

Backup and Disaster Recovery replication can fail if one of the destination KMSs fail

Even with KMS HA enabled on the destination cluster, BDR replication can still fail if one of the KMSs cannot be reached.

Workaround: Ensure all of the KMSs are online. Then, perform the replication again.

Affected Versions: Cloudera Manager 5.15.x, 5.16.x

Cloudera Manager installation fails on MariaDB 10.2.8 and later

When installing Cloudera Manager using MariaDB 10.2.8 or later, the Cloudera Manager web server doesn't come up and the install process ends with a failed status. The `cloudera-scm-server.log` includes the following SQL error:

```
2019-08-28 04:37:10,171 FATAL main:org.hsqldb.cmdline.SqlFile: SQL Error at 'UTF-8' line
57:
  "alter table ROLE_CONFIG_GROUPS
  drop column REVISION_ID"
  Key column 'REVISION_ID' doesn't exist in table
2019-08-28 04:37:10,171 FATAL main:org.hsqldb.cmdline.SqlFile: Rolling back SQL
transaction.
2019-08-28 04:37:10,172 ERROR main:com.cloudera.enterprise.dbutil.SqlFileRunner: Exception
while executing ddl scripts.
  com.mysql.jdbc.exceptions.jdbc4.MySQLSyntaxErrorException: Key column 'REVISION_ID'
doesn't exist in table
```

Note that MariaDB 10.2.8 is provided by default in some operating systems, including SLES 12 SP4.

Workaround: Replace the default MariaDB 10.2.x version with MariaDB 10.0.x.

Affected Versions: MariaDB 10.2.8 and later

Cloudera Issue: OPSAPS-52340

Failures during Cloudera Manager installation or upgrade

Cloudera Manager agent installation or upgrade fails due to misconfigured or problematic third-party repositories that were interfering with the process.

Cloudera Bug: OPSAPS-45576

Workaround: None

Affected versions: Cloudera Manager 5.14, 5.15

Fix versions: Cloudera Manager 5.16.1

Upgrade to CDH 5.15.1 fails with CM 5.15.1 and OpenJDK

Due to a Java version parsing issue with OpenJDK during Cloudera Manager upgrade, upgrades where OpenJDK is deployed fail.



Note: OpenJDK is only supported in Cloudera Manager and CDH 5.16.

Cloudera Bug: OPSAPS-47620

Fixed Version: Cloudera Manager and CDH 5.16.1, 5.15.2

After upgrade to Cloudera Manager 6, Kafka brokers and MirrorMaker marked as having stale configurations

The configurations that will be marked as stale are: - **BROKER_JAVA_OPTS** of the Kafka Broker role, and **MM_JAVA_OPTS** of the Kafka MirrorMaker role.

Workaround: Add the following options to the **Additional Broker Java Options** and **Additional MirrorMaker Java Options** Kafka configuration properties:

- `-Dcom.sun.management.jmxremote.host=127.0.0.1`
- `-Dcom.sun.management.jmxremote.local.only`

Restart the Kafka brokers.

Affected Versions: Cloudera Manager 5.15.0

Fixed Versions: Cloudera Manager 5.16.1, 5.15.2

Cloudera Bug: OPSAPS-47365

Parcel distribution fails and times out on Ubuntu 14.04 with Cloudera Express license

Parcel distribution to hosts from Cloudera Manager server is known to fail and timeout on Ubuntu 14.04 with a Cloudera Express license.

Workaround: None

Affected Versions: Cloudera Manager 5.15.0 and higher

Cloudera Issue: OPSAPS-44590

Flume service with Kafka over TLS marked as stale after upgrading Cloudera Manager from 5.15.0

After upgrading from Cloudera Manager 5.15.0, if you are using Flume + with Kafka over TLS, your Flume service will be marked as stale. To resolve this issue, restart the Flume service after the upgrade.

Affected Versions: Cloudera Manager 5.15.0

Flume uses default values for Kafka TLS settings

Cloudera Manager 5.15 added support for automatically configuring TLS settings when connecting to a Kafka service. In 5.15.0, the Flume configuration generated by Cloudera Manager when Flume is configured to connect to a secure Kafka cluster contains the default truststore location and password. Kafka TLS configuration inheritance does not handle non-default truststore file and password, which causes the connection to fail when the default values have been changed.

Workaround: Edit the Flume configuration file in Cloudera Manager to override the default parameters and use empty values for the Kafka producer and consumer truststore settings by adding the following lines to the Flume configuration file:

For Kafka channels:

```
AGENT_NAME.channels.CHANNEL_NAME.kafka.producer.ssl.truststore.location=  
AGENT_NAME.channels.CHANNEL_NAME.kafka.producer.ssl.truststore.password=  
AGENT_NAME.channels.CHANNEL_NAME.kafka.consumer.ssl.truststore.location=  
AGENT_NAME.channels.CHANNEL_NAME.kafka.consumer.ssl.truststore.password=
```

For Kafka sources:

```
AGENT_NAME.sources.SOURCE_NAME.kafka.consumer.ssl.truststore.location=  
AGENT_NAME.sources.SOURCE_NAME.kafka.consumer.ssl.truststore.password=
```

For Kafka sinks:

```
AGENT_NAME.sinks.SINK_NAME.kafka.producer.ssl.truststore.location=  
AGENT_NAME.sinks.SINK_NAME.kafka.producer.ssl.truststore.password=
```

AGENT_NAME, CHANNEL_NAME, SOURCE_NAME and SINK_NAME need to be replaced with the actual agent, channel, source and sink names respectively.

Affected Versions: Cloudera Manager 5.15.0

Fixed Versions: Cloudera Manager 5.15.1, 5.15.2, 5.16.1

After upgrading from Cloudera Manager 5.15.0, if you are using Flume with Kafka over TLS, your Flume service will be marked as stale. To resolve the problem, restart the Flume service after the upgrade.

Cloudera Issue: OPSAPS-45669

Upgrades from Cloudera Enterprise 5.15 to 6.0 Not Supported

You cannot upgrade from Cloudera Manager or CDH 5.15 to the upcoming release of Cloudera Manager and CDH 6.0.

Cloudera Manager read-only user can access sensitive cluster information

Due to a security vulnerability, a Cloudera Manager read-only user can access sensitive cluster information.

Products affected: Cloudera Manager

Releases affected:

- Cloudera Manager 5.12 and all prior releases
- Cloudera Manager 5.13.0, 5.13.1, 5.13.2, 5.13.3
- Cloudera Manager 5.14.0, 5.14.1, 5.14.2, 5.14.3

Users affected: All

Date of detection: April 18th, 2018

Detected by: Cloudera

Severity: High

Impact: Sensitive Information Disclosure

CVE: CVE-2018-10815

Immediate action required: Upgrade Cloudera Manager to a release with the fix.

Addressed in release/refresh/patch: Cloudera Manager 5.15.0 and higher, 5.14.4

Knowledge base: For the latest update on this issue, see the corresponding Knowledge article: TSB 2018-306: Cloudera Manager Information Disclosure

Accumulo health check fails on Cloudera Manager upgrade

After upgrading Cloudera Manager from 5.14.0 to 5.15.0, a health check failed with the message "ACCUMULO16-ACCUMULO16_MASTER-1 has undesired health: BAD"

Workaround:

Products affected: Cloudera Manager, Accumulo

Cloudera Bug: CDH-69677

Cloudera Manager should not omit LDAP GROUP_MAPPING passwords from NodeManager

After upgrading Cloudera Manager, clusters on CDH 5.5 or higher will be marked as stale if they had set a value for the **Hadoop User Group Mapping LDAP TLS/SSL Truststore Password** or **Hadoop User Group Mapping LDAP Bind User Password** configuration properties.

Workaround: Restart the stale services to pick up this fix. The configurations that will be marked stale are:

```
hadoop.security.group.mapping.ldap.ssl.keystore.password,  
hadoop.security.group.mapping.ldap.bind.password.
```

Products affected: Cloudera Manager

Affected versions: Cloudera Manager 5.5 or higher

Fixed versions: Cloudera Manager 5.16.1

Cloudera Bug: OPSAPS-45440

Import step in Hive replication can fail in 5.15.0 due to a Hive concurrency issue

In Cloudera Manager 5.15.0, BDR uses multi-threaded import and export for Hive replication by default. Because of a known Hive concurrency issue, the import step may occasionally fail with a `NullPointerException` for either the add partition or create table operations.

Workaround: Most of the time, running the replication again will be successful. If you continue to run into this issue, you can disable multi-threaded import and export by setting the **Number of concurrent HMS connections** on the **Advanced** tab of the replication schedule to 0 or fewer.

Products affected: Cloudera Manager

Affected versions: Cloudera Manager 5.15.0 and higher

Cloudera Bug: CDH-67557

Hard Restart of Cloudera Manager Agents May Cause Subsequent Service Errors

If a “hard restart” or “hard stop” operation is performed on a Cloudera Manager Agent, the restarted agent will erroneously restart roles that existed prior to the restart and, subsequently, 60 days later, these roles may experience errors or be killed.

Products affected: Cloudera Manager

Releases affected: All versions of Cloudera Manager 5.x

Cloudera Bug: OPSAPS-43550, TSB-308

Knowledge base: For the latest update on this issue, see the corresponding Knowledge article: [TSB 2018-308: Hard Restart of Cloudera Manager Agents May Cause Subsequent Service Errors](#)

Editing Cloudera Manager configurations in Google Chrome may lead to unexpected configuration loss

Google Chrome version 66, released on 28 April 2018, introduced behavior changes that can trigger unintended modification of existing configuration values in Cloudera Manager. If a page contains a password field, the browser aggressively applies auto-complete form data to the first text field and first password form fields. These unintended changes to the fields could be accidentally submitted unless the user pays careful attention to the visual cues (auto-filled fields that are highlighted) or change summary view presented just before submitting the changes.

Products affected: Cloudera Manager

Releases affected: All releases of Cloudera Manager

For the latest update on this issue see the corresponding Knowledge article: [TSB 2018-314: Editing Cloudera Manager configuration in Google Chrome may lead to unexpected configuration loss](#).

Impala Dynamic Resource Pools wrongly gives everyone access to root pool (and all child pools)

When you use ACLs for YARN queue management in the Cloudera Manager Admin Console and you do not specify any users or groups for the root pool, all users are allowed to access the pools.

Products affected: Cloudera Manager Admin Console: Dynamic Resource Pools

Releases affected: Cloudera Manager 5.12.1 or higher

Workaround: Specify a user or group for the root pool.

Cloudera Bug: OPSAPS-45046

Cross-site scripting vulnerability in Cloudera Manager

Several pages in the Cloudera Manager Admin console are vulnerable to a cross-site scripting attack.

Products affected: Cloudera Manager Admin Console

Releases affected:

- Cloudera Manager releases lower than 5.12
- Cloudera Manager 5.12.0, 5.12.1, 5.12.2
- Cloudera Manager 5.13.0, 5.13.1
- Cloudera Manager 5.14.0, 5.14.1

Users roles affected: Cluster Administrator, Full Administrator

Date of detection: January 19, 2018

Detected by: Shafeeque Olaseery Kunnikkal of Ingram Micro Asia Ltd, Security Consulting & Services

Severity (Low/Medium/High): High

Impact: A cross-site scripting vulnerability can be used by an attacker to perform malicious actions. One probable form of attack is to steal the credentials of a victim's Cloudera Manager account.

CVE: CVE-2018-5798

Immediate action required: Upgrade to a release in which this issue has been fixed.

Addressed in release/refresh/patch:

- Cloudera Manager 5.13.2 and higher
- Cloudera Manager 5.14.2 and higher

Knowledge base: For the latest update on this issue, see the corresponding Knowledge article: [XSS Scripting Vulnerability in Cloudera Manager](#)

Backup and Disaster Recovery (BDR) HDFS and Hive Replications may fail when replicating from source clusters which use HDFS HA and are managed by Cloudera Manager 5.13.3 or 5.14.2 to destination clusters managed by Cloudera 5.14.2.

All HDFS and Hive replication schedules fail when all the following conditions are met:

- The destination cluster is managed by Cloudera Manager 5.14.2.
- The source cluster is managed by Cloudera Manager 5.13.3+ or 5.14.2+.
- The source cluster uses HDFS HA.

Customers may observe logs with an error exception "java.net.UnknownHostException" that causes all replication jobs on to fail.

There is a workaround described below, which can have a performance impact on replication jobs (especially if the environment was previously running Cloudera Manager 5.13 or 5.14). The preferred solution is to upgrade to a later release. Upgrading is strongly recommended for any customer running HDFS or Hive replications and using Cloudera Manager 5.14.2, and is mandatory when any of the source cluster replicating data to the 5.14.2 clusters is a cluster managed by Cloudera Manager 5.13.3+ or 5.14.2+.

Products affected: Cloudera Manager Backup and Disaster Recovery

Releases affected: Cloudera Manager 5.14.2 (when used as the destination cluster of HDFS and/or Hive replication)

Users affected: Customers using HDFS or Hive Replication

Severity (Low/Medium/High): High

Root Cause and Impact:

In HDFS Replication, Cloudera Manager first runs a process on the source cluster that lists the files to be copied. Due to a bug, inconsistent arguments are issued to this process. This results in an attempt to copy an incorrect file that can't be accessed, resulting in an exception: `java.net.UnknownHostException`.

This issue affects all HDFS & Hive replication schedules when the destination cluster is managed by Cloudera Manager 5.14.2 and the source cluster is managed by either Cloudera Manager 5.13.3+ or 5.14.2+.

Immediate action required: If you use BDR, do not upgrade a destination environment to Cloudera Manager 5.14.2. Upgrade to Cloudera Manager 5.14.3 or higher when it becomes available.

If you have already upgraded to Cloudera Manager 5.14.2, you can workaround this bug by disabling the process that runs the file listing (copy listing) on the source cluster. This results in file listing phase happening in the destination environment, which can be slower when there is significant network latency between environments. This feature became available in Cloudera Manager 5.13.0.

To disable the feature that runs the copy listing phase, set the feature flag to false with the Cloudera Manager API as follows:

API endpoint for setting feature flag (HTTP PUT call):

```
api/v18/cm/config
```

Body of the API call:

```
{"items":[{"name":"feature_flag_run_copylist_source","value":"false"}]}
```

Please note that if you are upgrading from Cloudera Manager 5.12 or earlier to Cloudera Manager 5.14.2, disabling this feature will not have any impact on runtime. If you upgraded from Cloudera Manager 5.13.0 or Cloudera Manager 5.14.x to 5.14.2, you will see some performance degradation in the file listing phase of HDFS & Hive Replication as the process will be running on the destination cluster. The impact is driven by the number of HDFS files and folders in the source cluster.

Addressed in release/refresh/patch: Cloudera Manager 5.14.3 and higher

Knowledge base: For the latest update on this issue, see the corresponding Knowledge article: [TSB 2018-305: Backup and Disaster Recovery \(BDR\) HDFS and Hive Replications may fail after upgrading to Cloudera Manager 5.14.2](#)

Hive Replications Can Fail Intermittently

Hive replication jobs can fail intermittently during the first step of replication while exporting Hive metadata. When this happens, the following message displays in Cloudera Manager: **The remote command failed with error message: Hive Replication Export Step failed.** The likelihood that this error can occur increases with the length of time it takes to export Hive metadata. If you have a very large Hive deployment (those containing many tables and partitions), and/or environments in which the Hive metastore server is under-resourced, you are more likely to see this issue. However, having many tables and partitions or having an under-resourced Hive metastore server do not cause this issue. The problem is caused by a bug and the only solution is to upgrade to a later release. Cloudera strongly recommends upgrading to anyone running Hive Replication and using Cloudera Manager 5.13.0 or 5.13.1 versions.

Affected Products: Cloudera Manager

Affected Versions: Cloudera Manager 5.13.0, 5.13.1

Who is Affected: Anyone using Hive Replication

Severity (Low/Medium/High): High

Cause and Impact: A file is created during the Hive Export Phase, and is used by later phases, such as the Transfer/Import phases. Because of a bug, sometimes the file is overwritten by another process. Hive Replication can thus fail intermittently because of `FileNotFoundException`s. A retry may or may not resolve the issue. This issue can potentially affect some or all Hive replication schedules.

Workaround: Upgrade to Cloudera Manager 5.13.2, 5.14.1, or higher versions

Addressed in release/refresh/patch: Cloudera Manager 5.13.2, 5.14.1, or higher versions.

For the latest updates on this issue, see the following corresponding Knowledge Base article:

[TSB 2018-276: Hive Replications can fail intermittently in Cloudera Manager 5.13.0, 5.13.1 versions](#)

HDFS DataNode stale configuration

After upgrading to Cloudera Manager 5.10 or higher, the following configuration for HDFS DataNodes will be marked as stale: `dfs.datanode.balance.max.concurrent.moves`.

The staleness is caused by the following new feature: HDFS balancer can now be configured to specify which hosts are included and excluded or which hosts are used as sources for transferring replicas. Additional properties for tuning the performance of the balancer can now also be configured starting with CDH 5.10.0.

Workaround: You can safely ignore this warning and defer restarting.

Affected versions: Cloudera Manager 5.10.0 and lower

Cloudera Bug: OPSAPS-36642

YARN MapReduce Job History stale configuration

When upgrading to Cloudera Manager 5.10, YARN will be marked as having stale configuration due to `mapreduce.jobhistory.loadedjob.tasks.max`. Unless you change this parameter and want the non-default value to take effect (only takes effect in CDH 5.9+), you can simply ignore this staleness and defer restarting YARN.

Cloudera bug: OPSAPS-32132

Logging issue slows down Backup and Disaster Recovery Hive and HDFS Replication jobs

When using Cloudera Manager for cross-cluster replication of HDFS and Hive, unnecessary warning logs are printed during the initial phase of replication ('copy listing'). These logs are printed for each and every file that Cloudera Backup and Disaster Recovery (BDR) copies. This causes a very verbose log output which slows down this phase of the replication job and can affect overall replication times considerably. This happens only when both source and target clusters are running Cloudera Manager 5.14.0.

Workaround: Upgrade to Cloudera Manager 5.14.1 (or higher), or:

1. In the Cloudera Manager Admin Console for the **source** cluster, go to the HDFS service and click **Configuration > HDFS Replication Environment Advanced Configuration Snippet (Safety Valve)**.
2. Add the following parameter to the configuration snippet: `HADOOP_ROOT_LOGGER=ERROR`
3. Save the Configuration.



Warning: With this change, the amount of logging available for replication jobs will be significantly reduced, which can make it difficult to debug any issues with BDR replication. We highly recommend you revert this configuration change after upgrading to a version that includes the fix (5.14.1 or higher).

Affected versions: Cloudera Manager 5.14.0

Fixed version: Cloudera Manager 5.15.0, Cloudera Manager 5.14.1

Cloudera bug: OPSAPS-44160

See TSB-289.

Cloudera Manager upgrade workflow incorrectly requires deploying some optional management roles

When a user upgrades to Cloudera Manager 5.14.0, after all the agents have been updated to 5.14.0, the user sees an Upgrade Wizard, and is asked to select where to place the Navigator Metadata Server, Navigator Audit Server, and the Reports Manager roles. The **Continue** button is disabled until user adds one of the roles.

Workaround:

1. When prompted to add Navigator and Reports Manager roles, click the main Cloudera Manager Logo to exit the wizard.
2. Note that this action skips over a review of the Java heap size of the Host Monitor and Service Monitor roles. (Normally, Cloudera Manager recommends a value if it finds the existing values to be too small for the number of nodes being managed). If configured properly, this is not a concern. For more information on resource requirements for Cloudera Manager, see: [Cloudera Manager Hardware Requirements](#).
3. To complete the upgrade, start (or restart) the Cloudera Management Service (not the Cloudera Manager server). From the Cloudera Manager **Status** page, click the drop-down list next to **Cloudera Management Service** and select **Start** or **Restart**.

Alternatively, when prompted to add the Navigator and Reports Manager roles, users can:

1. Add the Navigator Metadata Server role and click **Continue**.
2. Configure a database for the Navigator Metadata Server and click **Continue**.
3. Stop and Delete the Navigator Metadata Server role from the **Cloudera Management Service > Instances** page immediately afterwards. See [Deleting Role Instances](#).

Affected Versions: Cloudera Manager 5.14.0

Fixed version: Cloudera Manager 5.14.1, 5.15.0

Cloudera Bug: OPSAPS-44629

See TSB-290

ADLS Credentials in Log Files

When you use Cloudera Manager to configure the ADLS Connector service using the **Less Secure** option for the **Credentials Protection Policy**, it is possible for Hive audit logs to include Microsoft Azure credentials. If you are using Navigator Audit Server, these credentials may appear in audit reports. To mitigate this problem, make sure that access to Hive logs is appropriately controlled and that Navigator users with **Auditing Viewer** roles are cleared to have access to the Hive credentials.

Affected Versions: Cloudera Manager 5.14.x

Cloudera Bug: CDH-56241

Installing Cloudera Manager on SLES Using the Installer Does Not Work

[Installing](#) Cloudera Manager 5 using the installer does not work on SLES due to PostgreSQL dependency issues

Workaround:

Perform a [regular installation](#).

Affected Versions: Cloudera Manager 5

Fixed Versions: Cloudera Manager 5.15.0

Cloudera Bug: OPSAPS-44284, OPSAPS-44628

Cloudera Navigator is not supported in multi-cluster deployments of Cloudera Manager versions 5.12.0, 5.12.1, and 5.13.0

Navigator is not able to uniquely identify more than one cluster in installations of Cloudera Manager that support multiple clusters. This problem only occurs for Cloudera Manager deployments of version 5.12.0, 5.12.1, or 5.13.0. Avoid multi-cluster installations with these releases if you expect to use Cloudera Navigator. This problem does not affect single cluster deployments or Altus clusters.

Fixed in versions: Cloudera Manager 5.14.0 (Cloudera Navigator 2.13.0), 5.13.1 (2.12.1), 5.12.2 (2.11.2)

Cloudera Bug: NAV-5615

Backup and Disaster Recovery Hive replication fails intermittently

In rare situations, due to a race condition, Hive Replication jobs can fail intermittently during the first step of replication with the following error: `The remote command failed with error message: Hive Replication Export Step failed.`

Workaround: Run the replication again or upgrade to a version that includes the bug fix.

Affected versions: Cloudera Manager 5.13.x

Fixed in: Cloudera Manager 5.14.0, 5.13.1

Cloudera bug: OPSAPS-43676

Impala logs missing from diagnostic bundle

Impala and Kudu role logs can be missing from the diagnostic bundles if their log directories have broken symlinks.

Affected versions: Cloudera Manager 5.14.0, 5.13.x

Cloudera bug: OPSAPS-41194

Accumulo logs are not included in diagnostic bundles

Accumulo logs are missing from the diagnostic bundle when collected by "date range".

Affected versions: Cloudera Manager 5.14.0, 5.13.x

Cloudera bug: OPSAPS-43206

Backup and Disaster Recovery replication fails between two clusters when data is copied from one encryption zone to another when using CDH 5.12.0 or higher.

This issue occurs when the two clusters use HDFS high availability, and the nameservice names are the same.

Workaround: Use unique nameservice names for the HDFS clusters.

Affected versions: CDH 5.12.0 and higher

Fixed in: Cloudera Manager 5.15.0, 5.14.3, 5.13.3

Cloudera bug: OPSAPS-43406, OPSAPS-43556

Backup and Disaster Recovery Hive replication fails during the Export Hive metastore step when using encrypted zones

If the `/user` or `/user/hdfs` directory is encrypted, then Hive replication fails with the following message: `User :hdfs not allowed to do 'DECRYPT_EEK' on some-key-name.`

Workaround: Configure a directory to store transient data during Hive replication:

1. Open the **Cloudera Manager Admin Console**.
2. Click **Backup > Replication Schedules**.
3. Select the replication schedule that fails and click **Actions > Edit Configuration**.
4. On the **Advanced** tab, specify an unencrypted directory for the **Export Path** or **Directory for metadata file** field. The field name depends on your Cloudera Manager version.

Affected versions: Cloudera Manager 5.13, 5.12.x, 5.11.x, 5.10.x, 5.9.x, 5.8.x

Cloudera bug: OPSAPS-43108

Backup and Disaster Recovery replication fails if the previous run is aborted

BDR replication fails with "another command" error if the previous run is aborted during the second or third step.

Workaround: Restart the Cloudera Manager Server. To prevent this error from occurring, do not abort replication during the second or third step.

Affected versions: Cloudera Manager 5.13

Fixed in: Cloudera Manager 5.14.0, 5.13.1

Backup and Disaster Recovery (Hive) replication fails if the user directory is encrypted

BDR replication fails if the user directory of the user specified in the **Run as Peer User** or **Run as username** field in the replication schedule is encrypted due to a change in how replications are performed. If no **Run as Peer User** is configured, then replication fails if the `hdfs` user directory is encrypted. For example, Hive replication fails on clusters where the `/user/testuser` directory is encrypted and the **Run as username** field specifies the "testuser" user.

Cloudera Manager 5 Release Notes

Affected versions: Cloudera Manager 5.13.0, 5.12.x

Fixed in: Cloudera Manager 5.14.0, 5.13.1

Workaround: Make the user directory unencrypted or configure a user for the replication who does not have an encrypted user directory. Alternatively, disable copylisting on source with the following Cloudera Manager API:

- API endpoint: `/cm/config`
- Payload: `{"items":[{"name":"feature_flag_run_copylist_source","value":"false"}]}`

Cloudera bug: OPSAPS-42445

Upgrade to CDH 5.13 or higher Requires Pre-installation of Spark 2.1 or Spark 2.2

If your cluster has Spark 2.0 or Spark 2.1 installed and you want to upgrade to CDH 5.13 or higher, you must first upgrade to Spark 2.1 release 2 or later before upgrading CDH. To install these versions of Spark, do the following before running the CDH Upgrade Wizard:

1. Install the Custom Service Descriptor (CSD) file. See

- [Installing Spark 2.1](#)
- [Installing Spark 2.2](#)



Note:

Spark 2.2 requires that JDK 1.8 be deployed throughout the cluster. JDK 1.7 is not supported for Spark 2.2.

See [Step 2: Install Java Development Kit](#).

2. Download, distribute, and activate the Parcel for the version of Spark that you are installing:

- **Spark 2.1 release 2:** The parcel name includes "cloudera2" in its name.
- **Spark 2.2 release 1:** The parcel name includes "cloudera1" in its name.

See [Managing Parcels](#).

Affected versions: CDH 5.13.0 and higher

Cloudera bug: CDH-56775

Record Service no longer supported as of CDH 5.13; upgrade fails after from some lower versions

When upgrading a cluster with Record Service installed from CDH 5.10.0 or lower to CDH 5.13 or higher, Record Service fails with a `java.lang.NoClassDefFoundError`. Record Service is no longer supported in CDH.

Affected CDH Versions: 5.10.0 and lower when upgrading to CDH 5.13.0 or higher.

Cloudera bug: CDH-59372

Sentry may require increased Java heap settings before upgrading CDH to 5.13

Before upgrading to CDH 5.13 or higher, you may need to increase the size of the Java heap for Sentry. A warning will be displayed during upgrade, but it's the user's responsibility to ensure this setting is adjusted properly before proceeding. See [Performance Guidelines](#).

Affected versions: CDH 5.13 or higher

Cloudera bug: OPSAPS-42541

Renaming an empty directory on Amazon S3 when using DynamoDB with S3Guard may fail

Renaming an empty directory on Amazon S3 when a DynamoDB table is used for S3Guard may fail and result in an error.

Affected versions: Cloudera Manager 5.12.0, 5.11.2, 5.11.1, 5.11.0

Fixed in: Cloudera Manager 5.12.0

Cloudera bug: CDH-51869

Apache bug: HADOOP-14036

Change in Hue Load Balancer version causes Hue Load Balancer and Server to be marked with stale configuration

After upgrading Cloudera Manager, the Hue Load Balancer and Server roles will be marked as having a stale configuration if all of these are true:

- The **Enable TLS/SSL for Hue** property is set to true.
- The Hue load balancer is enabled

Workaround: If your cluster uses Apache httpd 2.4 as the Hue load balancer, restart the Hue service promptly. If your cluster uses an earlier version of httpd, there is no urgency to restart the Hue service. (Apache httpd 2.4 is installed automatically by some recent versions of Linux, or may have been explicitly installed.)

Cloudera bug: OPSAPS-40700, OPSAPS-41850

Hue Load Balancer SSL Handshake error

Fixed an issue with the SSL handshake in Hue when Apache httpd 2.4 or higher is used as the load balancer. The Hue load balancer previously set the `ProxyPreserveHost` directive to `On`, when it should have been set to `Off`. This causes problems making SSL connections when using Apache httpd 2.4 or higher. The error caused problems when verifying the CN, which older versions of Apache httpd did not encounter because they did not properly verify the CN.

When upgrading Cloudera Manager, the Hue load balancer may be marked as having a stale configuration. If you are experiencing issues connecting to Hue with SSL, restart the Hue service to update the configuration.

Cloudera bug: OPSAPS-40700

Apache MapReduce Jobs May Fail During Rolling Upgrade to CDH 5.11.0 or CDH 5.11.1

In CDH 5.11, Cloudera introduced four new counters that are reported by MapReduce jobs. During a rolling upgrade from a cluster running CDH 5.10.x or lower to CDH 5.11.0 or CDH 5.11.1, a MapReduce job with an application master running on a host running CDH 5.10.x or lower may launch a map or reduce task on one of the newly-upgraded CDH 5.11.0 or CDH 5.11.1 hosts. The new task will attempt to report the new counter values, which the old application master will not understand, causing an error in the logs similar to the following:

```
2017-06-08 17:43:37,173 WARN [Socket Reader #1 for port 41187]
org.apache.hadoop.ipc.Server: Unable to read call parameters for client 10.17.242.22 on
connection protocol org.apache.hadoop.mapred.TaskUmbilicalProtocol for rpcKind
RPC_WRITABLE
java.lang.ArrayIndexOutOfBoundsException: 23
    at
    ...
```

This error could cause the task and the job to fail.

Workaround:

Avoid performing a rolling upgrade to CDH 5.11.0 or CDH 5.11.1 from CDH 5.10.x or lower. Instead, skip CDH 5.11.0 and CDH 5.11.1 if you are performing a rolling upgrade, and upgrade to CDH 5.12 or higher, or CDH 5.11.2 or higher when the release becomes available.

Cloudera bug: DOCS-2384, TSB-241

Rolling Upgrade with GPL Extras Parcel Causes Oozie to Fail

After performing a rolling upgrade where the GPL Extras parcel is upgraded, you must restart the Oozie service after completing the upgrade to let it pick up the latest client configurations. Otherwise, jobs newly submitted through Oozie may fail.

Affected versions: 5.14.0, 5.13.x, 5.12.x, 5.11.x

Cloudera bug: OPSAPS-41564

Maintenance State Minimal Block Replication staleness after upgrade

Upgrading to Cloudera Manager 5.12 or later may show `Maintenance State Minimal Block Replication` as a stale configuration under HDFS, suggesting a restart. It is safe to ignore this warning and delay restart.

Cloudera bug: OPSAPS-39102

Affected versions: Cloudera Manager 5.12 or higher

YARN ACL configuration property staleness after upgrade

After upgrading to Cloudera Manager 5.12, the following YARN configuration properties may show staleness warnings:

- ACL for viewing a job - `mapreduce.job.acl-view-job`
- ACL for modifying a job - `mapreduce.job.acl-modify-job`
- Enable MapReduce ACLs - `mapreduce.cluster.acls.enabled`

It is safe to defer restart if you are not using YARN job view/modify ACLs.

Affected versions: Cloudera Manager 5.12 or higher

Cloudera bug: OPSAPS-33586

Adding a service to an existing cluster

When you try to add a service such as Impala, Solr, Key-Value Store, or Sentry to an existing cluster, a timeout occurs.

Workaround

Deactivate the following parcels before you try to add a service: `KEYTRUSTEE`, `SQOOP_TERADATA_CONNECTOR` and any custom parcels.

Fixed in Cloudera Manager versions 5.8.5, 5.9.2, 5.10.1, and 5.11.

Cloudera bug: OPSAPS-38304

Services die due to HDFS taking too long to start

On Cloudera Manager managed clusters where HDFS takes a long time to come up after a restart, some dependent services may fail to start.

Workaround:

You can either manually start cluster services while waiting between steps or configure HDFS clients to automatically retry.

To manually start the cluster services, perform the following steps:

1. Start ZooKeeper and, if it's used, Kudu.
2. Start HDFS and wait for the HDFS NameNode to finish starting.
3. Run the following command:`hdfs dfsadmin -safemode get`.

This command turns safe mode off.

4. Start the remaining services.

To configure HDFS clients to automatically retry, perform the following steps:

1. In the Cloudera Manager web UI, select the HDFS service.
2. Select **Configuration > Gateway**.
3. In the **HDFS Client Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml** field, add the following property: `dfs.client.retry.policy.enabled`.
4. Set the value to true.
5. Add a description.
6. Save the changes.
7. Select **Actions > Deploy Client Configuration**.
8. Start the cluster.

Cloudera bug: CDH-54889

hostname parameter is not passed to Impala catalog role

[IMPALA-5253](#) contained a security fix for clusters using Impala with TLS (SSL) security enabled. This fix was also made in several maintenance versions of CDH that require you to upgrade Cloudera Manager. If you upgrade to a CDH version with this fix without upgrading Cloudera Manager, Impala will not function when TLS is enabled for Impala. You should upgrade Cloudera Manager first if you want to move to a CDH version with the security fix.

This issue affects upgrades of Cloudera Manager and CDH to version 5.11.1.

There are two ways you can workaround this issue:

- Upgrade to one of the following versions of Cloudera Manager before upgrading CDH:
 - 5.13.0
 - 5.12.1
 - 5.11.2
 - 5.10.2
 - 5.9.3
 - 5.8.5

-or-

- *Before upgrading CDH*, set the `-hostname` option to the fully-qualified domain name of the Catalog Server using the **Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)** configuration property:

```
-hostname=fully-qualified-domain-name of Impala Catalog Server
```

(To set this property, in Cloudera Manager, go to the Impala service, select the **Configuration** tab and search for the property.

CDH versions with the Impala security fix:

- 5.11.1
- 5.10.2
- 5.9.3
- 5.8.5

Cloudera bug: OPSAPS-41218

Zookeeper Package Installation fails with Debian 7

Zookeeper installation fails because Debian 7 installs its own version of the Zookeeper package instead of the Cloudera version. As versions change, this may also affect additional Cloudera packages.

Workaround:

- For manual installations of Zookeeper, run the following `apt-get` command instead of the documented command when installing Zookeeper:

```
apt-get install zookeeper=3.4.5+cdh5.10.2+108-1.cdh5.10.2.p0.4~wheezy-cdh5.10.2
```

- To ensure the installation uses the current Cloudera version of all packages, create a file called `cloudera.pref` in the `/etc/apt/preferences.d` directory and add the following content before running the `apt-get` command to manually install CDH or using Cloudera Manager to install the packages:

```
Package: *  
Pin: release o=Cloudera, l=Cloudera  
Pin-Priority: 501
```

This change gives Cloudera packages higher priority than Debian-supplied packages.

The change assumes your base packages have the priority set to '500'. You can check whether the priority is correctly set by running the `apt-cache policy package_name` command. The output should list Cloudera packages as highest priority. This step must be done before both manual and Cloudera Manager-assisted package installations.

Cloudera bug: DOCS-2278

Automated Cloudera Manager installer fails on Ubuntu 16.04

Running the `cloudera-manager-installer.bin` installer file (as described in the [documentation](#)) fails on Ubuntu 16.04 LTS (Xenial).

Affected versions: Cloudera Manager 5.11.0

Fixed in: Cloudera Manager 5.11.1 and higher.

Cloudera bug: DOCS-2037

Reboot of HDFS nodes running RHEL 7 or CentOS 7 shortly after enabling High Availability or migrating NameNodes can lead to data loss or other issues

This issue affects Cloudera Manager 5.7.0 and 5.7.1 installed on RHEL 7.x or CentOS 7.x.

On RHEL 7.x and CentOS 7.x systems, certain configuration actions intended to be executed only once during enablement of HDFS HA or migration of HDFS roles might be erroneously re-executed if a system shutdown or reboot occurs within next 24 hours of either operation. This can result in data loss and inconsistent state if the user has performed a NameNode format or JournalNode format, which are part of both the enablement of HDFS HA (adding a standby NameNode) and relocation (migration) of NameNodes roles to other hosts. Edits to HDFS metadata stored in JournalNodes could be lost in these situations, requiring manual repair to recover recent changes to HDFS data. If you experience this issue, please contact Cloudera Support for assistance.

For the latest update on this issue see the corresponding Knowledge article: [TSB 2017-161: Reboot of HDFS nodes running RHEL 7/CentOS 7 after enabling High Availability or migrating NameNodes can lead to data loss or other issues.](#)

Resolution:

Upgrade to Cloudera Manager 5.7.2 or higher immediately, and perform a hard restart of Cloudera Manager agents.



Note: A hard agent restart will cause CDH roles to be restarted but can be performed after the upgrade in a scheduled, rolling manner, leveraging the HA capabilities of the CDH services. A normal rolling restart from the Cloudera Manager UI will not perform a hard restart.

Fixed in: Cloudera Manager 5.7.2 and higher.

The `Limit Nonsecure Container Executor Users` property has no effect in Cloudera 5.11

Workaround:

To configure the `yarn.nodemanager.linux-container-executor.nonsecure-mode.limit-users` property to true or false, you can do so using the YARN parameter **NodeManager Advanced Configuration Snippet (Safety Valve) for yarn-site.xml**.

Cloudera bug: OPSAPS-27702

Graceful shutdown of Kafka brokers does not work as expected

In CM 5.11.0, the new **Graceful Shutdown Timeout** configuration property does not work as expected. As a result, Kafka takes an additional 30 seconds (by default) to shut down, but will still only have 30 seconds to complete its controlled shutdown, before Cloudera Manager forcibly shuts down Kafka brokers regardless of the configured timeout.

Workaround:

Wait longer for the shutdown to occur, or reduce the **Graceful Shutdown Timeout** to a very low value, for example: 1. (You should restore the original value when you upgrade to a Cloudera Manager release with the fix for this issue, although it will not cause a critical problem if left at a low value.) There is no workaround that allows you to increase the real time that Kafka has to perform shutdown until this issue is fixed.

Fixed in Cloudera Manager 5.11.1.

Cloudera bug: OPSAPS-40106

Spark Gateway roles should be added to every host

If you are using Cloudera Manager 5.9.0, 5.9.1, or 5.10.0 and have hosts with the NodeManager role but without the Spark Gateway role, you must [add the Spark Gateway role](#) to all NodeManager hosts and [redploy the client configurations](#). If you do not use this workaround, Cloudera Manager fails to locate the `topology.py` file, which can cause task localization issues and failures for very large jobs.

Fixed in Cloudera Manager 5.10.1.

Cloudera bug: OPSAPS-39119

HBase configuration is supplied for Hive when HBase service is not selected

Cloudera Manager provides configuration for the `hive-site.xml` file even if the HBase Service setting in Hive is not selected. This can cause unnecessary errors when Hive-on-Spark attempts to connect to HBase.

Workaround:

1. In Cloudera Manager, go to the Hive service.
2. Click the **Configuration** tab.
3. Search for the following Advanced Configuration Snippets:
 - **HiveServer2 Advanced Configuration Snippet (Safety Valve) for hive-site.xml**
 - **Hive Client Advanced Configuration Snippet (Safety Valve) for hive-site.xml**
4. Click the **+** icon to add the following property and value to both Advanced Configuration Snippets:
Name: `spark.yarn.security.tokens.hbase.enabled`
Value: `false`
5. Click **Save Changes**.
6. Click the stale configuration icon to restart stale services, including any dependent services. Follow the on-screen prompts.
7. Click **Actions > Deploy Client Configuration**.

8. When you execute a query that uses a HBase-backed table, set this parameter back to `true`. You can do this by changing the configuration as described above (set each value to `true`) or by using a `set` command and then issuing the query. For example:

```
set spark.yarn.security.tokens.hbase.enabled=true;
SELECT * FROM HBASE_BACKED_TABLE ...
```

Fixed in Cloudera Manager 5.8.5 and 5.9.3.

Cloudera bug: OPSAPS-39021

Hive Replication fails when Impala is SSL enabled but Hadoop services are not

Support for Impala Replication with SSL enabled was added in Cloudera Manager 5.7.4, 5.8.2, 5.9.0 and higher.

Workaround:

1. In Cloudera Manager, go to the HDFS service.
2. Click the **Configuration** tab.
3. Search for and enter the appropriate values for the following parameters:
 - **Cluster-Wide Default TLS/SSL Client Truststore Location**
 - **Cluster-Wide Default TLS/SSL Client Truststore Password**

Impala metadata replication is not supported, but will be supported in a later maintenance release.

This support also ensures that connection to Impala is successful when SSL is enabled even if Kerberos is not enabled.

Cloudera bug: OPSAPS-38700, OPSAPS-38720

Fixed in Cloudera Manager 5.11, 5.10.1, 5.9.2, 5.8.4, 5.7.6.

python-psycopg2 Dependency

Cloudera Manager 5.8 and higher has a new dependency on the package `python-psycopg2`. This package is not available in standard SLES 11 and SLES 12 repositories. You need to add this repository or install it manually to any machine that runs the Cloudera Manager Agent before you install or upgrade Cloudera Manager.

If the Cloudera Manager Server and Agent run on the same host, install the Cloudera Manager Server first and then add the `python-psycopg2` repository or package. After adding the repository or package, install the Cloudera Manager Agent.

Workaround

Download the `python-psycopg2` repository or package from the following URL by selecting the correct SLES version: <https://software.opensuse.org/package/python-psycopg2>.

Pauses in Cloudera Manager after adding peer

Pauses and slow performance of the Cloudera Manager Admin Console occur after creating a peer.

Workaround:

Increase the memory allocated to Cloudera Manager:

1. On the Cloudera Manager server host, edit the `/etc/default/cloudera-scm-server` file.
2. In the `Java Options` section, change the value of the `-Xmx` argument of the `CMF_JAVA_OPTS` property from `-Xmx2G` to `-Xmx4G`. For example:

```
export CMF_JAVA_OPTS="-Xmx4G -XX:MaxPermSize=256m -XX:+HeapDumpOnOutOfMemoryError -XX:HeapDumpPath=/tmp"
```

3. Restart the Cloudera Manager server:

```
sudo service cloudera-scm-server restart
```

Fixed in Cloudera Manager 5.11

Cloudera bug: OPSAPS-38868

Cannot select time period in custom charts in C5.9

The quick time selection (30m, 1h, and so on) on custom dashboards does not work in 5.9.x.

Fixed in Cloudera Manager 5.9.1 and higher.

Cloudera bug: OPSAPS-37190

Cloudera Manager 5.8.2 allows you to select nonexistent CDH 5.8.1 package installation

The Cloudera manager 5.8.2 install/upgrade wizard allows you to select CDH 5.8.1 as a package installation, even though CDH 5.8.1 does not exist. The installation fails with an error message similar to the following:

```
[Errno 14] HTTPS Error 404 - Not Found
```

Workaround

Return to the package selection page, and select **Latest Release of CDH5 component compatible with this version of Cloudera Manager** or **CDH 5.8.2**.

Cloudera bug: OPSAPS-36860

Error when distributing parcels: No such torrent

Parcel distribution might fail with an error message conforming to:

```
Error when distributing to <host>: No such torrent: <parcel_name>.torrent
```

Workaround

Remove the file `/opt/cloudera/parcel-cache/<parcel_name>.torrent` from the host.

Cloudera bug: OPSAPS-37183

Hive Replication Metadata Transfer Step fails with Temporary AWS Credential Provider

Hive Replication Schedules that use Amazon S3 as the **Source** or **Destination** fail when using temporary AWS credentials. The following error displays:

```
Message: Hive Replication Metadata Transfer Step Failed -
com.cloudera.com.amazonaws.services.s3.model.AmazonS3Exception:
Status Code: 403, AWS Service: Amazon S3, AWS Request ID: 76D1F6A02792908A,
AWS Error Code: null, AWS Error Message: Forbidden,
S3 Extended Request ID:
Xy3nAS4HSPKLA6hHKvpqReBud7M1Fhk7On0HttYGE0eKPHKwiFkTPQxEVU82OZq5d8omSrdhbcI=.
```

Cloudera bug: OPSAPS-37514

Hive table Views do not get restored from S3

When creating a Hive Replication schedule that copies Hive data from S3 and you select the **Reference Data From Cloud** option, Hive table Views are not restored correctly and result in a Null Pointer Exception when querying data from the view.

Fixed in Cloudera Manager 5.9.1 and higher.

Cloudera bug: OPSAPS-37549

ACLs are not replicated when restoring Hive data from S3

ACLs are never replicated when the **Enable Access Control Lists** option in the configuration of the HDFS service is not selected the first time a Replication Schedule that replicates from S3 to Hive runs. Enabling it and re-running the restore operation does not restore the ACLs.

Cloudera bug: OPSAPS-37004

Snapshot diff is not working for Hive to S3 replication when data is deleted on source

If you have enabled snapshots on an HDFS folder and a Hive table uses an external file in that folder, and then you replicate that data to S3 and delete the file on the source cluster, the file is not deleted in subsequent replications to S3, even if the **Delete Permanently** option is selected.

Cloudera bug: OPSAPS-36910

Block agents from heartbeating to a Cloudera Manager with different UUID until agent restart

If for some reason Cloudera Manager server's identity (Cloudera Manager guid) is changed, agents drop heartbeat requests and will not follow any requested commands from Cloudera Manager server. As a result, the agents report bad health. This situation can be fixed by taking either of the following approaches:

- Restore the previous Cloudera Manager server guid

OR

- Remove the `cm_guid` file from each of the agents and then restart the agent.

Cloudera bug: OPSAPS-34847

Cloudera Manager set catalogd default jvm memory to 4G can cause out of memory error on upgrade to Cloudera Manager 5.7 or higher

After upgrading to 5.7 or higher, you might see a reduced Java heap maximum on Impala Catalog Server due to a change in its default value. Upgrading from Cloudera Manager lower than 5.7 to Cloudera Manager 5.8.2 no longer causes any effective change in the Impala Catalog Server Java Heap size.

When upgrading from Cloudera Manager 5.7 or later to Cloudera Manager 5.8.2, if the Impala Catalog Server Java Heap Size is set at the default (4GB), it is automatically changed to either 1/4 of the physical RAM on that host, or 32GB, whichever is lower. This can result in a higher or a lower heap, which could cause additional resource contention or out of memory errors, respectively.

Cloudera bug: OPSAPS-34039

Cloudera Manager 5.7.4 installer does not show Key Trustee KMS

A fresh install of Cloudera Manager tries to install Key Trustee KMS 5.8.2 when trying to install the latest version. You must either choose 5.7.0 as the Key Trustee KMS version, or manually provide a link to the 5.7.4 bits.

Class Not Found Error when upgrading to Cloudera Manager 5.7.2

When you upgrade to version 5.7.2 of Cloudera Manager, the client configuration for all services is marked *stale*.

Workaround:

From the **Cluster** menu, select **Deploy Client Configuration** to redeploy the client configuration.

Kerberos setup fails on Debian 8.2

This issue is due to the following Debian bug: <https://bugs.debian.org/cgi-bin/bugreport.cgi?bug=777579;msg=5;att=0>.

Workaround:

1. Log in to the host where the Cloudera Manager server is running.
2. Edit the `systemd/system/krb5-admin-server.service` file and add `/etc/krb5kdc` to the `ReadWriteDirectories` section.
3. Run the following commands:

```
systemctl daemon-reload
sudo service krb5-admin-server restart
```

4. Generate the credentials.

Cloudera bug: OPSAPS-33713

Password in Cloudera Manager's `db.properties` file is not redacted

The `db.properties` file is managed by customers and is populated manually when the Cloudera Manager Server database is being set up for the first time. Since this occurs before the Cloudera Manager Server has even started, encrypting the contents of this file is a completely different challenge as compared to that of redacting configuration files.

Releases affected: 5.3 and higher

Cloudera bug: OPSAPS-24813

Cluster provisioning fails

In some cases, provisioning of a cluster may fail at the start of the process. This does not happen in all cases and is mainly noticed on RHEL 6 and especially when some hosts are reporting bad health.

Releases affected: 5.5.0-5.5.3, 5.6.0-5.6.1, 5.7.0

Releases containing the fix: 5.5.4, 5.7.1

For releases containing the fix, parcel activation and first run command now completes as expected, even when some hosts report bad health.

This issue is fixed in Cloudera Manager 5.5.4 and 5.7.1 and higher.

Cloudera bug: OPSAPS-33564

Cloudera Manager can run out of memory if a remote repository URL is unreachable

If one of the URLs specified in on the **Parcel Settings** page (**Hosts > Parcels > Configuration**) becomes unreachable, Cloudera Manager may run out of memory.

Workaround:

Do one of the following:

- If the URL is incorrect, enter the correct URL.
- Deselect the **Automatically Download New Parcels** setting on the **Parcel Settings** page.
- Set the value of the **Parcel Update Frequency** on the **Parcel Settings** to a large interval such as several days.

Cloudera bug: OPSAPS-31732

Clients can run Hive on Spark jobs even if Hive dependency on Spark is not configured

In CDH 5.7 and higher, when Hive and Spark on YARN both are configured, but Hive is not configured to depend on Spark on YARN, clients can set the execution engine to `spark` and Hive on Spark jobs will still be executed but will run in an unsupported mode. These jobs may not appear in the Spark History Server.

Workaround: Configure Hive to depend on Spark on YARN.

Cloudera bug: OPSAPS-32023

The YARN NodeManager connectivity health test does not work for CDH 5

The NodeManager connectivity is always GOOD (green) even if the ResourceManager considers the NodeManager to be LOST or DECOMMISSIONED.

Cloudera bug: OPSAPS-31251

Workaround: None.

HDFS HA clusters see NameNode failures when KDC connectivity is bad

When KDC connectivity is bad, the JVM takes 30 seconds before retrying or declaring failure to connect. Meanwhile, the JournalNode write timeout (which needs KDC authentication for the first write, or under troubled connectivity), is only 20 seconds.

Workaround: In `krb5.conf`, set the `kdc_timeout` parameter value to 3 seconds. In Cloudera Manager, perform the following steps:

1. Go to **Administration > Settings > Kerberos**.
2. Add the `kdc_timeout` parameter to the **Advanced Configuration Snippet (Safety Valve) for [libdefaults]** section of `krb5.conf` property. This should give the JVM enough time to try connecting to a KDC before the JournalNode timeout.

The HDFS File browser in Cloudera Manager fails when HDFS federation is enabled

Workaround: Use the command-line `hdfs dfs` commands to directly manipulate HDFS files when federation is enabled. CDH supports HDFS federation.

Hive Metastore canary fails to drop database

The Hive Metastore canary fails to drop the database due to [HIVE-11418](#).

Cloudera bug: OPSAPS-27632

Workaround: Turn off the Hive Metastore canary by disabling the Hive Metastore Canary Health Test:

1. Go to the Hive service.
2. Click the **Configuration** tab.
3. Select **SCOPE > Hive Metastore Server**.
4. Select **CATEGORY > Monitoring**.
5. Deselect the **Hive Metastore Canary Health Test** checkbox for the Hive Metastore Server Default Group.
6. Click **Save Changes** to commit the changes.

Cloudera Manager upgrade fails due to incorrect Sqoop 2 path

Sqoop 2 will not start or loses data when upgrading from Cloudera Manager 5.4.0 or 5.4.1 to Cloudera Manager 5.4.3, or when upgrading Cloudera Manager 3 or 4 to Cloudera Manager 5.4.0 or 5.4.1. This is due to Cloudera Manager erroneously configuring the Derby path with "reposito~~y~~" instead of "repository". To upgrade Cloudera Manager, use one of the following workarounds:

- **Workaround for Upgrading from Cloudera Manager 3 or 4 to Cloudera Manager 5.4.0 or 5.4.1**
 1. Log in to your Sqoop 2 server host using SSH and move the Derby database files to the new location, usually from `/var/lib/sqoop2/repositoy` to `/var/lib/sqoop2/repository`.
 2. Start Sqoop2. If you found this problem while upgrading CDH, run the Sqoop 2 database upgrade command using the **Actions** drop-down menu for Sqoop 2.
- **Workaround for Upgrading from Cloudera Manager 5.4.0 or 5.4.1 to Cloudera Manager 5.4.3**
 1. Log in to your Sqoop 2 server host using SSH and move the Derby database files to the new location, usually from `/var/lib/sqoop2/repositoy` to `/var/lib/sqoop2/repository`.

2. Start Sqoop2, or if you found this problem while upgrading CDH, run the Sqoop 2 database upgrade command using the **Actions** drop-down menu for Sqoop 2.

NameNode incorrectly reports missing blocks during rolling upgrade

During a rolling upgrade to any of the CDH releases listed below, the NameNode may report missing blocks after rolling back multiple DataNodes. This is caused by a race condition with block reporting between the DataNode and the NameNode. No permanent data loss occurs, but data can be unavailable for up to six hours before the problem corrects itself.

Releases affected: CDH 5.0.6, 5.1.5, 5.2.5, 5.3.3, 5.4.1, 5.4.2.

Releases containing the fix:: CDH 5.2.6, 5.3.4, 5.4.3

Workaround:

- **To avoid the problem** - Cloudera advises skipping the affected releases and installing a release containing the fix. For example, do not upgrade to CDH 5.4.2; upgrade to CDH 5.4.3 instead.
- **If you have already completed an upgrade to an affected release, or are installing a new cluster** - You can continue to run the release, or upgrade to a release that is not affected.

Cloudera bug: OPSAPS-27205

Using ext3 for server dirs easily hit inode limit

Using the ext3 filesystem for the Cloudera Manager command storage directory may exceed the maximum subdirectory size of 32000.

Workaround: Either decrease the value of the **Command Eviction Age** property so that the directories are more aggressively cleaned up, or migrate to the ext4 filesystem.

Cloudera bug: OPSAPS-26951

Backup and disaster recovery replication does not set MapReduce Java options

Replication used for backup and disaster recovery relies on system-wide MapReduce memory options, and you cannot configure the options using the Advanced Configuration Snippet.

Cloudera bug: OPSAPS-25503

Kafka 1.2 CSD conflicts with CSD included in Cloudera Manager

If the Kafka CSD was installed in Cloudera Manager to 5.3 or lower, the old version must be uninstalled, otherwise it will conflict with the version of the Kafka CSD bundled with Cloudera Manager 5.4.

Workaround: Remove the Kafka 1.2 CSD before upgrading Cloudera Manager to 5.4:

1. Determine the location of the CSD directory:
 - a. Select **Administration > Settings**.
 - b. Click the **Custom Service Descriptors** category.
 - c. Retrieve the directory from the **Local Descriptor Repository Path** property.
2. Delete the Kafka CSD from the directory.

Cloudera bug: OPSAPS-25995

Recommission host does not deploy client configurations

The failure to deploy client configurations can result in client configuration pointing to the wrong locations, which can cause errors such as the NodeManager failing to start with "Failed to initialize container executor".

Workaround: Deploy client configurations first and then restart roles on the recommissioned host.

Cloudera bug: OPSAPS-25995

Hive on Spark is not supported in Cloudera Manager and CDH 5.4 and CDH 5.5

You can configure Hive on Spark, but it is not recommended for production clusters.

Cloudera bug: OPSAPS-25983

Upgrade wizard incorrectly upgrades the Sentry DB

There's no Sentry DB upgrade in 5.4, but the upgrade wizard says there is. Performing the upgrade command is not harmful, and taking the backup is also not harmful, but the steps are unnecessary.

Cloudera bug: OPSAPS-25405

Cloudera Manager does not correctly generate client configurations for services deployed using CSDs

HiveServer2 requires a Spark on YARN gateway on the same host in order for Hive on Spark to work. You must deploy Spark client configurations whenever there's a change in order for HiveServer2 to pick up the change.

CSDs that depend on Spark will get incomplete Spark client configuration. Note that Cloudera Manager does not ship with any such CSDs by default.

Cloudera bug: OPSAPS-25695

Workaround: Use `/etc/spark/conf` for Spark configuration, and ensure there is a Spark on YARN gateway on that host.

Solr, Oozie and HttpFS fail when KMS and TLS/SSL are enabled using self-signed certificates

When the KMS service is added and TLS/SSL is enabled, Solr, Oozie and HttpFS are not automatically configured to trust the KMS's self-signed certificate and you might see the following error.

```
org.apache.oozie.service.AuthorizationException: E0501: Could not perform authorization operation,
sun.security.validator.ValidatorException: PKIX path building failed:
sun.security.provider.certpath.SunCertPathBuilderException:
unable to find valid certification path to requested target
```

Cloudera bug: OPSAPS-23392, CDH-23460, CDH-23189

Workaround: You must explicitly load the relevant truststores with the KMS certificate to allow these services to communicate with the KMS. To do so, edit the truststore location and password for Solr, Oozie and HttpFS (found under the HDFS service) as follows.

1. Go to the Cloudera Manager Admin Console.
2. Go to the Solr/Oozie/HDFS service.
3. Click the **Configuration** tab.
4. Search for "<service> TLS/SSL Certificate Trust Store File" and set this property to the location of truststore file.
5. Search for "<service> TLS/SSL Certificate Trust Store Password" and set this property to the password of the truststore.
6. Click **Save Changes** to commit the changes.

Cloudera Manager 5.3.1 upgrade fails if Spark standalone and Kerberos are configured

CDH upgrade fails if Kerberos is enabled and Spark standalone is installed. Spark standalone does not work in a kerberized cluster.

Cloudera bug: OPSAPS-24983

Workaround: To upgrade, remove the Spark standalone service first and then proceed with upgrade.

Adding Key Trustee KMS 5.4 to Cloudera Manager 5.5 displays warning

Adding the Key Trustee KMS service to a CDH 5.4 cluster managed by Cloudera Manager 5.5 displays the following message, even if Key Trustee KMS is installed:

"The following selected services cannot be used due to missing components: keytrustee-keyprovider. Are you sure you wish to continue with them?"

Workaround: Verify that the Key Trustee KMS parcel or package is installed and click **OK** to continue adding the service.

KMS and Key Trustee ACLs do not work in Cloudera Manager

ACLs configured for the KMS (File) and KMS (Navigator Key Trustee) services do not work since these services do not receive the values for `hadoop.security.group.mapping` and related group mapping configuration properties.

Cloudera bug: OPSAPS-24483

Workaround:

KMS (File): Add all configuration properties starting with `hadoop.security.group.mapping` from the NameNode `core-site.xml` to the KMS (File) property, **Key Management Server Advanced Configuration Snippet (Safety Valve) for core-site.xml**

KMS (Navigator Key Trustee): Add all configuration properties starting with `hadoop.security.group.mapping` from the NameNode `core-site.xml` to the KMS (Navigator Key Trustee) property, **Key Management Server Proxy Advanced Configuration Snippet (Safety Valve) for core-site.xml**.

Exporting and importing Hue database sometimes times out after 90 seconds

Executing 'dump database' or 'load database' of Hue from Cloudera Manager returns "command aborted because of exception: Command timed-out after 90 seconds". The Hue database can be exported to JSON from within Cloudera Manager. Unfortunately, sometimes the Hue database is quite large and the export times out after 90 seconds.

Cloudera bug: OPSAPS-24470

Workaround: Ignore the timeout. The command should eventually succeed even though Cloudera Manager reports that it timed out.

Changing the Key Trustee Server hostname requires editing keytrustee.conf

If you change the hostname of your active or passive Key Trustee Server, you must edit the `keytrustee.conf` file. This issue typically arises if you replace an active or passive server with a server having a different hostname. If the same hostname is used on the replacement server, there are no issues.

Cloudera bug: OPSAPS-24133

Workaround: Use the same hostname on the replacement server.

Hosts with Impala Llama roles must also have at least one YARN role

When integrated resource management is enabled for Impala, host(s) where the Impala Llama role(s) are running must have at least one YARN role. This is because Llama requires the `topology.py` script from the YARN configuration. If this requirement is not met, you may see errors such as:

```
"Exception running /etc/hadoop/conf.cloudera.yarn/topology.py
java.io.IOException: Cannot run program "/etc/hadoop/conf.cloudera.yarn/topology.py"
```

in the Llama role logs, and Impala queries may fail.

Cloudera bug: OPSAPS-23728

Workaround: Add a YARN gateway role to each Llama host that does not already have at least one YARN role (of any type).

The high availability wizard does not verify that there is a running ZooKeeper service

If one of the following is true:

- 1. ZooKeeper present and not running and the HDFS dependency on ZooKeeper dependency is not set
- 2. ZooKeeper absent

the enable high-availability wizard fails.

Cloudera bug: OPSAPS-23709

Workaround: Before enabling high availability, do the following:

1. Create and start a ZooKeeper service if one does not exist.
2. Go to the HDFS service.
3. Click the **Configuration** tab.
4. Select **Scope > Service-Wide**
5. Set the **ZooKeeper Service** property to the ZooKeeper service.
6. Click **Save Changes** to commit the changes.

Cloudera Manager Installation Path A fails on RHEL 5.7 due to PostgreSQL conflict

On RHEL 5.7, `cloudera-manager-installer.bin` fails due to a PostgreSQL conflict if PostgreSQL 8.1 is already installed on your host.

Cloudera bug: OPSAPS-23392

Workaround: Remove PostgreSQL from host and rerun `cloudera-manager-installer.bin`.

Spurious warning on Accumulo 1.6 gateway hosts

When using the Accumulo shell on a host with only an Accumulo 1.6 Service gateway role, users will receive a warning about failing to create the directory `/var/log/accumulo`. The shell works normally otherwise.

Workaround: The warning is safe to ignore.

Cloudera bug: OPSAPS-21699

Accumulo 1.6 service log aggregation and search does not work

Cloudera Manager log aggregation and search features are incompatible with the log formatting needed by the Accumulo Monitor. Attempting to use either the "Log Search" diagnostics feature or the log file link off of an individual service role's summary page will result in empty search results.

Cloudera bug: OPSAPS-21675

Severity: High

Workaround: Operators can use the Accumulo Monitor to see recent severe log messages. They can see recent log messages below the WARNING level via a given role's process page and can inspect full logs on individual hosts by looking in `/var/log/accumulo`.

Cloudera Manager incorrectly sizes Accumulo Tablet Server max heap size after 1.4.4-cdh4.5.0 to 1.6.0-cdh4.6.0 upgrade

Because the upgrade path from Accumulo 1.4.4-cdh4.5.0 to 1.6.0-cdh4.6.0 involves having both services installed simultaneously, Cloudera Manager will be under the impression that worker hosts in the cluster are oversubscribed on memory and attempt to downsize the max heap size allowed for 1.6.0-cdh4.6.0 Tablet Servers.

Cloudera bug: OPSAPS-21806

Severity: High

Workaround: Manually verify that the Accumulo 1.6.0-cdh4.6.0 Tablet Server max heap size is large enough for your needs. Cloudera recommends you set this value to the sum of 1.4.4-cdh4.5.0 Tablet Server and Logger heap sizes.

Accumulo installations using LZO do not indicate dependence on the GPL Extras parcel

Accumulo 1.6 installations that use LZO compression functionality do not indicate that LZO depends on the GPL Extras parcel. When Accumulo is configured to use LZO, Cloudera Manager has no way to track that the Accumulo service now relies on the GPL Extras parcel. This prevents Cloudera Manager from warning administrators before they remove the parcel while Accumulo still requires it for proper operation.

Cloudera bug: OPSAPS-21680

Workaround: Check your Accumulo 1.6 service for the configuration changes mentioned in the Cloudera documentation for using Accumulo with CDH prior to removing the GPL Extras parcel. If the parcel is mistakenly removed, reinstall it and restart the Accumulo 1.6 service.

Created pools are not preserved when Dynamic Resource Pools page is used to configure YARN or Impala

Pools created on demand are not preserved when changes are made using the Dynamic Resource Pools page. If the Dynamic Resource Pools page is used to configure YARN or Impala services in a cluster, it is possible to specify pool placement rules that create a pool if one does not already exist. If changes are made to the configuration using this page, pools created as a result of such rules are not preserved across the configuration change.

Cloudera bug: OPSAPS-19942

Workaround: Submit the YARN application or Impala query as before, and the pool will be created on demand once again.

User should be prompted to add the AMON role when adding MapReduce to a CDH 5 cluster

When the MapReduce service is added to a CDH 5 cluster, the user is not asked to add the AMON role. Then, an error displays when the user tries to view MapReduce activities.

Cloudera bug: OPSAPS-20379

Workaround: Manually add the AMON role after adding the MapReduce service.

Enterprise license expiration alert not displayed until Cloudera Manager Server is restarted

When an enterprise license expires, the expiration notification banner is not displayed until the Cloudera Manager Server has been restarted. The enterprise features of Cloudera Manager are not affected by an expired license.

Cloudera bug: OPSAPS-15711

Workaround: None.

Configurations for decommissioned roles not migrated from MapReduce to YARN

When the **Import MapReduce Configuration** wizard is used to import MapReduce configurations to YARN, decommissioned roles in the MapReduce service do not cause the corresponding imported roles to be marked as decommissioned in YARN.

Cloudera bug: OPSAPS-16670

Workaround: Delete or decommission the roles in YARN after running the import.

The HDFS command Roll Edits does not work in the UI when HDFS is federated

The HDFS command Roll Edits does not work in the Cloudera Manager UI when HDFS is federated because the command does not know which nameservice to use.

Cloudera bug: OPSAPS-10034

Workaround: Use the API, not the Cloudera Manager UI, to execute the Roll Edits command.

Cloudera Manager reports a confusing version number if you have oozie-client, but not oozie installed on a CDH 4.4 node

In CDH versions before 4.4, the metadata identifying Oozie was placed in the client, rather than the server package. Consequently, if the client package is not installed, but the server is, Cloudera Manager will report Oozie has been present but as coming from CDH 3 instead of CDH 4.

Cloudera bug: OPSAPS-15778

Workaround: Either install the oozie-client package, or upgrade to at least CDH 4.4. Parcel based installations are unaffected.

Cloudera Manager does not work with CDH 5.0.0 Beta 1

When you upgrade from Cloudera Manager 5.0.0 Beta 1 with CDH 5.0.0 Beta 1 to Cloudera Manager 5.0.0 Beta 2, Cloudera Manager won't work with CDH 5.0.0 Beta 1 and there's no notification of that fact.

Cloudera bug: OPSAPS-17802

Workaround: None. Do a new installation of CDH 5.0.0 Beta 2.

On CDH 4.1 secure clusters managed by Cloudera Manager 4.8.1 and higher, the Impala Catalog server needs advanced configuration snippet update

Impala queries fail on CDH 4.1 when Hive "Bypass Hive Metastore Server" option is selected.

Workaround: Add the following to Impala catalog server advanced configuration snippet for `hive-site.xml`, replacing `Hive_Metastore_Server_Host` with the host name of your Hive Metastore Server:

```
<property>
<name>hive.metastore.local</name>
<value>false</value>
</property>
<property>
<name>hive.metastore.uris</name>
<value>thrift://Hive_Metastore_Server_Host:9083</value>
</property>
```

Rolling Upgrade to CDH 5 is not supported.

Rolling upgrade between CDH 4 and CDH 5 is not supported. Incompatibilities between major versions means rolling restarts are not possible. In addition, rolling upgrade will *not* be supported from CDH 5.0.0 Beta 1 to any later releases, and may not be supported between *any future beta versions* of CDH 5 and the General Availability release of CDH 5.

Workaround: None.

Error reading .zip file created with the Collect Diagnostic Data command.

After collecting Diagnostic Data and using the Download Diagnostic Data button to download the created zip file to the local system, the zip file cannot be opened using the FireFox browser on a Macintosh. This is because the zip file is created as a Zip64 file, and the unzip utility included with Macs does not support Zip64. The zip utility must be version 6.0 or later. You can determine the zip version with `unzip -v`.

Cloudera bug: OPSAPS-13850

Workaround: Update the unzip utility to a version that supports Zip64.

After JobTracker failover, complete jobs from the previous active JobTracker are not visible.

When a JobTracker failover occurs and a new JobTracker becomes active, the new JobTracker UI does not show the completed jobs from the previously active JobTracker (that is now the standby JobTracker). For these jobs the "Job Details" link does not work.

Cloudera bug: OPSAPS-13864

Severity: Med

Workaround: None.

After JobTracker failover, information about rerun jobs is not updated in Activity Monitor.

When a JobTracker failover occurs while there are running jobs, jobs are restarted by the new active JobTracker by default. For the restarted jobs the Activity Monitor will not update the following: 1) The start time of the restarted job will remain the start time of the original job. 2) Any Map or Reduce task that had finished before the failure happened will not be updated with information about the corresponding task that was rerun by the new active JobTracker.

Cloudera bug: OPSAPS-13879

Severity: Med

Workaround: None.

Installing on AWS, you must use private EC2 hostnames.

When installing on an AWS instance, and adding hosts using their public names, the installation will fail when the hosts fail to heartbeat.

Severity: Med

Workaround:

Use the Back button in the wizard to return to the original screen, where it prompts for a license.

Rerun the wizard, but choose "Use existing hosts" instead of searching for hosts. Now those hosts show up with their internal EC2 names.

Continue through the wizard and the installation should succeed.

If HDFS uses Quorum-based Storage without HA enabled, the SecondaryNameNode cannot checkpoint.

If HDFS is set up in non-HA mode, but with Quorum-based storage configured, the `dfs.namenode.edits.dir` is automatically configured to the Quorum-based Storage URI. However, the SecondaryNameNode cannot currently read the edits from a Quorum-based Storage URI, and will be unable to do a checkpoint.

Severity: Medium

Workaround: Add to the NameNode's advanced configuration snippet the `dfs.namenode.edits.dir` property with both the value of the Quorum-based Storage URI as well as a local directory, and restart the NameNode. For example,

```
<property> <name>dfs.namenode.edits.dir</name>
<value>qjournal://jn1HostName:8485;jn2HostName:8485;jn3HostName:8485/journalhdfs1,file:///dfs/edits</value>
</property>
```

Changing the rack configuration may temporarily cause mis-replicated blocks to be reported.

A rack re-configuration will cause HDFS to report mis-replicated blocks until HDFS rebalances the system, which may take some time. This is a normal side-effect of changing the configuration.

Severity: Low

Workaround: None

Cannot use '/' as a mount point with a Federated HDFS Nameservice.

A Federated HDFS Service does not support nested mount points, so it is impossible to mount anything at '/'. Because of this issue, the root directory will always be read-only, and any client application that requires a writeable root directory will fail.

Severity: Low

Workaround:

1. In the CDH 4 HDFS Service > Configuration tab of the Cloudera Manager Admin Console, search for "nameservice".
2. In the Mountpoints field, change the mount point from "/" to a list of mount points that are in the namespace that the Nameservice will manage. (You can enter this as a comma-separated list - for example, "/hbase, /tmp, /user" or by clicking the plus icon to add each mount point in its own field.) You can determine the list of mount points by running the command `hadoop fs -ls /` from the CLI on the NameNode host.

Historical disk usage reports do not work with federated HDFS.

Severity: Low

Workaround: None.

(CDH 4 only) Activity monitoring does not work on YARN activities.

Activity monitoring is not supported for YARN in CDH 4.

Severity: Low

Workaround: None

HDFS monitoring configuration applies to all Nameservices

The monitoring configurations at the HDFS level apply to all Nameservices. So, if there are two federated Nameservices, it's not possible to disable a check on one but not the other. Likewise, it's not possible to have different thresholds for the two Nameservices.

Severity: Low

Workaround: None

Supported and Unsupported Replication Scenarios and Limitations

See [Data Replication](#).

Restoring snapshot of a file to an empty directory does not overwrite the directory

Restoring the snapshot of an HDFS file to an HDFS path that is an empty HDFS directory (using the Restore As action) will result in the restored file present inside the HDFS directory instead of overwriting the empty HDFS directory.

Workaround: None.

HDFS Snapshot appears to fail if policy specifies duplicate directories.

In an HDFS snapshot policy, if a directory is specified more than once, the snapshot appears to fail with an error message on the Snapshot page. However, in the HDFS Browser, the snapshot is shown as having been created successfully.

Severity: Low

Workaround: Remove the duplicate directory specification from the policy.

Hive replication fails if "Force Overwrite" is not set.

The Force Overwrite option, if checked, forces overwriting data in the target metastore if there are incompatible changes detected. For example, if the target metastore was modified and a new partition was added to a table, this option would force deletion of that partition, overwriting the table with the version found on the source. If the Force Overwrite option is not set, recurring replications may fail.

Severity: Med

Workaround: Set the Force Overwrite option.

Cloudera Manager set cataloged default JVM memory to 4G can cause an out of memory error during upgrade to Cloudera Manager 5.7 and higher

When upgrading from Cloudera Manager 5.7 or higher to Cloudera Manager 5.8.2, if the Impala Catalog Server Java Heap Size is set at the default (4GB), it is automatically changed to either 1/4 of the physical RAM on that host, or 32GB, whichever is lower. This can result in a higher or a lower heap, which could cause additional resource contention or out of memory errors, respectively.

Host Monitor Cannot Monitor NTP Status on SLES Hosts

In Cloudera Manager 5.11.0 and lower, Host Monitor queries the NTP status using the `ntpd` command. Later versions of SLES disable this command by default.

Workaround: Add the following line to `/etc/ntp.conf`:

```
enable mode7
```

Affected versions: Cloudera Manager 5.11.0 and lower

Fixed versions: Cloudera Manager 5.11.1 and higher, Cloudera Manager 5.12.0 and higher

Cloudera Issue: OPSAPS-38268

Issues Fixed in Cloudera Manager 5

The following sections describe issues fixed in each Cloudera Manager 5 release.

Issues Fixed in Cloudera Manager 5.16.2

The API method "ApiHiveCloudReplicationArguments" is missing for Cloudera Manager 5.15

The Cloudera Manager Python API client released in 5.16.1 missed changes that were introduced in version 5.15, including changes related to Backup and Disaster Recovery (BDR) Hive Replication. Those missing changes have been ported.

Cloudera Bug: OPSAPS-48873

Backup and Disaster Recovery fails if `dfs.nameservices` is overridden with multiple nameservice names

This fix handles the scenario for multiple nameservices in the `dfs.nameservices` configuration. Now the feature cross-checks that with the `fs.defaultFS` configured in `core-site.xml`.

Cloudera Bug: OPSAPS-48579

Error handling differs between multi-thread and single-thread Hive replication

This fix aligns multi-thread Hive replication error handling with the single-thread case. Multi-thread Hive replication now has the appropriate error types, suppressed by the `replication.hive.ignoreTableNotFound` and `replication.hive.ignoreDataBaseNotFound` safety valves implicitly.

Cloudera Bug: OPSAPS-49987

Issues Fixed in Cloudera Manager 5.16.1



Note: The first release of Cloudera Manager 5.16 is numbered 5.16.1.

Yarn Application Masters fail when a concurrent Deploy Client Configuration job is executed

Yarn Application Masters fail because the `container-executor.cfg` file is missing when a YARN application is executed concurrently with a Deploy Client Configurations job. The container-executor hierarchy has been moved to `/var/lib/yarn-ce`, which is not modified by the Deploy Client Configuration command.

No changes are required for Cloudera Manager/CDH 5.16.0 and higher. Use the following **Workaround** for Cloudera Manager/CDH 5.15.1 and 5.14.4:

1. On the Cloudera Manager server host, add the following line to the `/etc/default/cloudera-scm-server` file:

```
export CMF_FF_YARN_SAFE_CONTAINER_EXECUTOR_DIR=true
```

2. Restart the Cloudera Manager Server:

```
service cloudera-scm-server restart
```

3. Restart the NodeManager(s) so they can pick up the new configuration change:
 - a. In the Cloudera Manager Admin console, go to the YARN service.
 - b. Click the **Instances** tab.
 - c. Select all hosts with the NodeManager **Role Type**.
 - d. Click **Actions for Selected** > **Restart**.

Cloudera Bug: OPSAPS-24398

Aborting a BDR replication does not stop and issues an empty error

The abort functionality now works as expected.

Cloudera Bug: OPSAPS-39694

HDFS replication fails when reserved raw data is replicated

Fixed an issue with HDFS replication where reserved raw data is replicated with the delete strategy.

Cloudera Bug: OPSAPS-47244

Threshold settings prevent correct configuration

When Threshold settings are saved with values such as Any or Never, a bug in the validation code prevented the user from saving the threshold to a specific value. This is now fixed.

Cloudera Bug: OPSAPS-46657

Gateway roles and Isilon replication jobs

Fixed a bug where Cloudera Manager did not allow Gateway roles for Isilon clusters, which prevented replications jobs from running.

Cloudera Bug: OPSAPS-46154

Save Changes remains enabled after clicking (possible autocomplete issue)

Google Chrome version 66, released on 28 April 2018, introduced behavior changes which can trigger unintended modification of existing configuration values in Cloudera Manager. If a page contains a password field, the browser aggressively applies auto-complete form data to the first text field and first password form fields. These changes are applied without user input or explicit notification (fields are highlighted).

Cloudera Bug: OPSAPS-46118

Spark cross-realm authentication fails

Spark now correctly respects `auth_to_local` name rules for HDFS services with cross-realm trust configured.

Cloudera Bug: OPSAPS-46103

Host installation or upgrade using key based authentication where private key is protected by passphrase does not work

In Cloudera Manager, host installation and upgrade does not work when using key-based host authentication with SSH where the private key is protected by a passphrase.

Cloudera Bug: OPSAPS-45571

Failures during Cloudera Manager installation or upgrade

Fixed an issue where Cloudera Manager agent installation or upgrade failed due to misconfigured or problematic third-party repositories that were interfering with the process.

Cloudera Bug: OPSAPS-45576

Cloudera Manager should not omit LDAP GROUP_MAPPING passwords from NodeManager

Some sensitive authentication configurations were not emitted for the Yarn Node Manager when using LDAP group mapping. This is now corrected for clusters running CDH 5.5 and higher where encryption is available for those sensitive configurations.

Cloudera Bug: OPSAPS-45440

Upgrading a license finishes on wrong page

The Enable Trial workflow previously finished on the upgrade page, now it goes back to the Home page upon completion.

Cloudera Bug: OPSAPS-45444

Flume uses default values for Kafka TLS settings

Fixed an issue where Flume used default values for Kafka TLS settings. For further details, see the original [Known Issue](#).

Note that this fix causes Flume to be marked as having a stale configuration. To resolve the staleness, restart Flume.

Cloudera Bug: OPSAPS-45669

Cloudera Manager upgrade page load issues in Internet Explorer

The new Cloudera Manager Upgrade page doesn't load in IE9, 10 or 11 due to using a feature that is not supported in IE. This is now fixed in 5.15.1.

Cloudera Bug: OPSAPS-45779

Do not expose configuration history to non-admin users

Users were previously able to go directly to the configuration history and rollback page using the URL, even though the link itself is not visible. This is now protected.

Cloudera Bug: OPSAPS-45781

Incorrect page reloads on Instances and All Hosts paged

The Instances page and the All Hosts page now do not reload when a command finishes.

Cloudera Bug: OPSAPS-45761

Kudu package missing from CDH installation using packages

Fixed an issue where Cloudera Manager did not install Kudu packages when CDH was installed using packages instead of parcels.

Cloudera Manager 5 Release Notes

Cloudera Bug: OPSAPS-45692

Reports Manager displays incorrectly

When the Reports Manager is not installed, Cloudera Manager Admin Console pages now display correctly and irrelevant links have been removed.

Cloudera Bug: OPSAPS-44873

Cloudera Manager fails when enabling Kerberos if TLS is already configured

When enabling kerberos for a cluster running TLS, the system cannot use the privileged ports (< 1024). Instead, the wizard will prompt user to use the appropriate port values.

Cloudera Bug: OPSAPS-33345

Kafka JMX connections listen on all interfaces unnecessarily

Kafka broker and MirrorMaker processes now listen on only the loopback interface for JMX connections. This caused Cloudera Manager to report that Kafka had stale configurations.

Cloudera Bug: OPSAPS-47365

BDR replicates Kudu HMS entries as-is

Hive tables created by Kudu are now excluded from Backup and Disaster Recovery replication.

Cloudera Bug: OPSAPS-46549

Restart warnings are incorrect after starting role with outdated configuration

Fixed an issue where some roles that required restarts were not correctly identified after starting a role marked as **Started with Outdated Configuration**.

Cloudera Bug: OPSAPS-45237

Host inspector should check system user group membership

The new Host Inspector will now warn if user doesn't belong to its own group

Cloudera Bug: OPSAPS-46592

Null Pointer Exception on Cloudera Manager Upgrade Page

Fixed a NPE on the new Cloudera Manager Upgrade page. This only happens when the agent could not retrieve the host status.

Cloudera Bug: OPSAPS-46616

Upgrade to CDH 5.15.1 fails with CM 5.15.1 and OpenJDK

Fixes Java version parsing issue with OpenJDK during Cloudera Manager upgrade.



Note: OpenJDK is only supported in Cloudera Manager and CDH 5.16.

Cloudera Bug: OPSAPS-47620

Issues Fixed in Cloudera Manager 5.15.2

Restart Warnings are incorrect after starting a role with an outdated configuration

Fixed an issue where some roles that required restarts were not correctly identified after starting a role marked as "Started with Outdated Configuration".

Cloudera Bug: OPSAPS-47775, OPSAPS-47493

Releases affected: 5.13, 5.14, 5.15

Upgrade to CDH 5.15.1 fails with CM 5.15.1 and OpenJDK

Fixes Java version parsing issue with OpenJDK during Cloudera Manager upgrade.



Note: OpenJDK is only supported in Cloudera Manager and CDH 5.16.

Cloudera Bug: OPSAPS-47620

Kafka JMX connections listen on all interfaces unnecessarily

Kafka broker and MirrorMaker processes now listen on only the loopback interface for JMX connections. This caused Cloudera Manager to report that Kafka had stale configurations.

Cloudera Bug: OPSAPS-47365

Flume uses default values for Kafka TLS settings

Fixed an issue where Flume used default values for Kafka TLS settings. For further details, see the original [Known Issue](#).

Note that this fix causes Flume to be marked as having a stale configuration. To resolve the staleness, restart Flume.

Cloudera Bug: OPSAPS-45669

Issues Fixed in Cloudera Manager 5.15.1

Cloudera Manager read-only user can access sensitive cluster information

Fixed an issue where, due to a security vulnerability, a Cloudera Manager read-only user can access sensitive cluster information.

For the latest update on this issue, see the corresponding Knowledge article:

TSB 2018-306: Cloudera Manager Information Disclosure

Cloudera Bug: OPSAPS-45688

Open Redirect and XSS in Cloudera Manager

Technical Service Bulletin 2018-321 (TSB)

One type of page in Cloudera Manager uses a `returnUrl` parameter to redirect the user to another page in Cloudera Manager once a wizard is completed. The validity of this parameter was not checked. As a result, the user could be automatically redirected to an attacker's external site or perform a malicious JavaScript function that results in cross-site scripting (XSS).

With this fix, Cloudera Manager no longer allows any value in the `returnUrl` parameter with patterns such as `http://`, `https://`, `///`, or `javascript`. The only exceptions to this rule are the SAML login/logout URLs, since they are explicitly configured and are not passed via the `returnUrl` parameter.

Products affected: Cloudera Manager

Releases affected:

- 5.15.0 and all earlier releases

Users affected: The following Cloudera Manager roles: "cluster administrator", "full administrators", and "configurators".

Date/time of detection: June 20, 2018

Detected by: Mohit Rawat & Ekta Mittal

Severity (Low/Medium/High): 8.8 High ([CVSS](#):3.0/AV:N/AC:L/PR:N/UI:R/S:U/C:H/I:H/A:H)

Impact: [Open redirects](#) can silently redirect a victim to an attacker's site. XSS vulnerabilities can be used to steal credentials or to perform arbitrary actions as the targeted user.

CVE: CVE-2018-15913

Immediate action required: Upgrade to Cloudera Manager 5.15.1 or higher

Addressed in release/refresh/patch:

- Cloudera Manager 5.15.1 and higher
- Cloudera Manager 6.0.0

Flume uses default values for Kafka TLS settings

Fixed an issue where Flume used default values for Kafka TLS settings. For further details, see the original [Known Issue](#).

Note that this fix causes Flume to be marked as having a stale configuration. To resolve the staleness, restart Flume.

Cloudera Bug: OPSAPS-45669

Abort Installation button causes errors

Abort Installation button causes errors When installing or upgrading hosts using the Upgrade or Installation wizards, when the ""abort"" link is clicked an error message is displayed. (waiting for answer in comments)

Cloudera Bug: OPSAPS-45655

Host Inspector reports identical JDK versions as different.

The Host Inspector was reporting JDKs installed with a different installation path as different JDK versions.

Cloudera Bug: OPSAPS-44931

Upgrading a license finishes on wrong page

The Enable Trial workflow previously finished on the upgrade page, now it goes back to the Home page upon completion.

Cloudera Bug: OPSAPS-45444

Cloudera Manager should not omit LDAP GROUP_MAPPING passwords from NodeManager

Some sensitive authentication configurations were not emitted for the Yarn Node Manager when using LDAP group mapping. This is now corrected for clusters running CDH 5.5 and higher where encryption is available for those sensitive configurations.

Cloudera Bug: OPSAPS-45440

Do not expose configuration history to non-admin users

Users were previously able to go directly to the configuration history and rollback page using the URL, even though the link itself is not visible. This is now protected.

Cloudera Bug: OPSAPS-45781

Cloudera Manager upgrade page load issues in Internet Explorer

The new Cloudera Manager Upgrade page doesn't load in IE9, 10 or 11 due to using a feature that is not supported in IE. This is now fixed in 5.15.1.

Cloudera Bug: OPSAPS-45779

Gateway roles and Isilon replication jobs

Fixed a bug where Cloudera Manager did not allow Gateway roles for Isilon clusters, which prevented replications jobs from running.

Cloudera Bug: OPSAPS-46154

Cloudera Backup and Disaster Recovery cannot import Hive metadata

Fixed a bug where Cloudera Backup and Disaster Recovery could not import Hive metadata that included a Partitioned View with a null StorageDescriptor.

Cloudera Bug: OPSAPS-45900

Threshold settings prevent correct configuration

When Threshold settings are saved with values such as Any or Never, a bug in the validation code prevented the user from saving the threshold to a specific value. This is now fixed.

Cloudera Bug: OPSAPS-46657

Incorrect page reloads on Instances and All Hosts paged

The Instances page and the All Hosts page now do not reload when a command finishes.

Cloudera Bug: OPSAPS-45761

Null Pointer Exception on Cloudera Manager Upgrade Page

Fixed a NPE on the new Cloudera Manager Upgrade page. This only happens when the agent could not retrieve the host status.

Cloudera Bug: OPSAPS-46616

Spark cross-realm authentication fails

Spark now correctly respects `auth_to_local` name rules for HDFS services with cross-realm trust configured.

Cloudera Bug: OPSAPS-46103

Yarn Application Masters fail when a concurrent Deploy Client Configuration job is executed

Yarn Application Masters fail because the `container-executor.cfg` file is missing when a YARN application is executed concurrently with a Deploy Client Configurations job. The container-executor hierarchy has been moved to `/var/lib/yarn-ce`, which is not modified by the Deploy Client Configuration command.

No changes are required for Cloudera Manager/CDH 5.16.0 and higher. Use the following **Workaround** for Cloudera Manager/CDH 5.15.1 and 5.14.4:

1. On the Cloudera Manager server host, add the following line to the `/etc/default/cloudera-scm-server` file:

```
export CMF_FF_YARN_SAFE_CONTAINER_EXECUTOR_DIR=true
```

2. Restart the Cloudera Manager Server:

```
service cloudera-scm-server restart
```

3. Restart the NodeManager(s) so they can pick up the new configuration change:
 - a. In the Cloudera Manager Admin console, go to the YARN service.
 - b. Click the **Instances** tab.
 - c. Select all hosts with the NodeManager **Role Type**.
 - d. Click **Actions for Selected > Restart**.

Cloudera Bug: OPSAPS-24398

Cloudera Manager does not list any CDH package options other than "custom"

When doing a package installation or upgrade of CDH, no versions of CDH5 are shown in the Cloudera Manager Admin Console.

Cloudera Bug: OPSAPS-45732

Issues Fixed in Cloudera Manager 5.15.0

[Cloudera Manager read-only user can access sensitive cluster information](#)

Fixed an issue where, due to a security vulnerability, a Cloudera Manager read-only user can access sensitive cluster information.

For the latest update on this issue, see the corresponding Knowledge article:

TSB 2018-306: Cloudera Manager Information Disclosure

Cloudera Bug: OPSAPS-45688

[Cross-site scripting vulnerability in Cloudera Manager](#)

Fixed the cross-site vulnerability described in [TSB-2018-285 \(XSS Scripting Vulnerability in Cloudera Manager\)](#). For more information, see [the Known Issue description](#)

Cloudera Bug: DOCS-3186

[Cannot install or start Cloudera Manager on SLES12](#)

Fixed an issue where [Path A](#) (non-production) installations of Cloudera Manager do not work on SLES version 12.

See TSB-279 for more information.

Cloudera Bug: OPSAPS-44284

[Extra entries in the Manage Host Override dialog](#)

Fixed an issue where duplicate hostnames appeared when editing configuration overrides for Host level configurations.

Cloudera Bug: OPSAPS-44814

[Cloudera logo disappears in Upgrade Wizard](#)

Fixed an issue where the Cloudera Manager logo might not get refreshed in the browser automatically when upgrading from Cloudera Manager 5.10 to 5.11 or higher. This issue caused the logo to disappear.

Cloudera Bug: OPSAPS-44676

[Errors thrown for Hive replications managed by a single Cloudera Manager](#)

Fixed an issue where an exception is thrown when collecting replication diagnostics for a Hive replication schedule between two clusters managed by a single Cloudera Manager.

Cloudera Bug: OPSAPS-45197

[YARN NodeManager stale due to missing CGroups](#)

Fixed an issue where the YARN NodeManager may show as being stale due to System Resources even when it is not. The `named-cpu` value may show as changed even though it was not modified.

Cloudera Bug: OPSAPS-43973

[Warning message "Cloud file creation time does not match saved time" in S3 replication](#)

Fixed multiple false warning log message issue during HDFS-S3 restore operation.

Cloudera Bug: OPSAPS-43681

[Concurrent Hive replications on different Hives fail with clashing command error](#)

Fixed an issue that caused Hive replications to fail due to conflicting Hive replication schedules when there are two clusters managed by the same Cloudera Manager.

Cloudera Bug: OPSAPS-43652

Impala Dynamic Resource Pools access

Fixed an issue where Cloudera Manager did not generate the `aclSubmitApps` element in the `fair-scheduler.xml` file when a user did not specify any user or groups for the root pool. Cloudera Manager now automatically generates `aclSubmitApps` with a space character. This change means the default is no access for all users.

Cloudera Bug: OPSAPS-45046

Cannot start CM5.13.1 after one-click manual install on SLES12 SP2

This fix upgrades the PostgreSQL JDBC driver. The new JDBC driver only supports PostgreSQL versions 8.2 and higher. RHEL5 defaults to PostgreSQL 8.1.

Use the following workaround for RHEL5

1. Download PostgreSQL 9.0-802 JDBC4: <https://jdbc.postgresql.org/download/postgresql-9.0-802.jdbc4.jar>
2. Copy it as `/usr/share/java/postgresql-connector-java.jar` and then install the server.

Cloudera Bug: OPSAPS-43624

AvroTypeException preventing Cloudera Manager Agent from connecting to the Cloudera Manager Server

Fixed an issue where an agent heartbeat fails and leads to the agent not connecting to the Cloudera Manager Server.

Cloudera Bug: OPSAPS-44971

Host Templates Regression on CM 5.13.x or 5.14

Fixed an issue where the Apply Host Template feature in the Add Host wizard was broken.

Cloudera Bug: OPSAPS-45334

Agent stops heartbeating due to Avro parsing errors.

Fixed an issue where, under certain conditions, the Cloudera Manager Agent stops heart beating with the Cloudera Manager Server due to Avro RPC parsing errors.

Cloudera Bug: OPSAPS-45282

Error "Mismatched input PATTERN expecting EOF" on the Detail Usage page of the Resource Manager

Fixed an issue where a user sees an error message about `Mismatched input PATTERN`.

Cloudera Bug: OPSAPS-42437

Agent exception "Could not evaluate resource {u'hard_limit': 500, u'soft_limit': 300}: 'unicode' object has no attribute 'items'"

Fixes an issue in Cloudera Manager agent where cgroup soft and hard memory limits were not applied on the host.

Cloudera Bug: OPSAPS-45229

Host installation or upgrade using key based authentication where private key is protected by passphrase does not work

In Cloudera Manager, host installation and upgrade does not work when using key-based host authentication with SSH where the private key is protected by a passphrase.

Cloudera Bug: OPSAPS-45571

Backup and Disaster Recovery replication fails between two clusters when data is copied from one encryption zone to another when using CDH 5.12.0 or higher

Fixed an issue where replication fails when two clusters use HDFS HA, and the nameservice names are the same.

Both the source and destination clusters must upgrade to a release that contains the fix.



Important: Cloudera recommends using unique names for the nameservices.

Cloudera Bug: OPSAPS-43556

Fixed in: Cloudera Manager 5.15.0, 5.14.3, 5.13.3

Logging issue slows down Hive and HDFS Replication jobs

Fixed an issue that occurred when using Cloudera Manager for cross-cluster replication of HDFS and Hive. Unnecessary warning logs are printed during the initial phase of replication ('copy listing'). These logs were printed for each and every file that Cloudera Backup and Disaster Recovery (BDR) copies. This causes very verbose log output which slows down this phase of the replication job and can affect overall replication times considerably. This happens only when both source and target clusters are running Cloudera Manager 5.14.0.

See TSB-290 for more information.

Fixed Versions: Cloudera Manager 5.15.0, 5.14.1

Cloudera Bug: OPSAPS-44160

Scheduled snapshot and replication jobs are skipped

Fixed an issue where scheduled Snapshot and Replication jobs can be skipped if there are too many commands scheduled to run at the same time (more than 100), or if Cloudera Manager has paused due to garbage collection.

Fixed Versions: Cloudera Manager 5.15.0, 5.14.2, 5.13.2

Cloudera Bug: OPSAPS-44153

Table statistics are lost when replicating an existing table without partitions

Cloudera Backup and Disaster Recovery added a workaround for the Hive bug HIVE-15653 and now preserves table statistics for unpartitioned tables during subsequent runs.

Fixed Versions: Cloudera Manager 5.15.0, 5.14.2, 5.13.2

Cloudera Bug: OPSAPS-44070

The Replication performance CSV file header does not match content

Fixed an issue where the header in the HDFS Replication performance file (in CSV format) does not match its contents.

Fixed in: Cloudera Manager 5.15.0, 5.14.2, 5.13.2

Cloudera Bug: OPSAPS-43997

BDR Cloud (s3) Replication job fail with delegation token creation issue

Fixed an issue where HDFS to Amazon S3 or Hive to Amazon S3 replication jobs fail when an expired delegation token is used to run a MapReduce/YARN job.

Cloudera Bug: OPSAPS-44008

Spark CSD should emit proper environment variables for older CDH version 5.4.x

Fixed configuration of Kerberos credentials for the Spark History Server in CDH 5.4.

Cloudera Bug: OPSAPS-43985

Hive Replication export step fails with StaleObjectStateException

Fixed an issue where replication commands sometimes fail because of concurrent executions in the database.

Cloudera Bug: OPSAPS-43964

Add "sentry" group in "hadoop.proxyuser.hive.groups" on upgrade to CDH 5.13.0 and higher

For CDH5.13.0 and higher, when Hive is used with Sentry, Hive Proxy Groups should contain either a wildcard (*) for the Sentry Group for Sentry to work with Hive. If this is not present, Cloudera Manager adds the Sentry Group to Hive Proxy Groups when upgrading CDH.

Cloudera Bug: OPSAPS-43809

Cloudera Manager Server reports 'Too many open files' error

Fixed an issue where collection of diagnostic bundles would lead to slow file descriptor leaks in Cloudera Manager Server that eventually cause the server to become non-operational or unresponsive.

Cloudera Bug: OPSAPS-44922

Cloudera Manager Upgrade (5.14.0 to 5.14.1) stuck at refreshing package metadata - configurator not found

Fixed a bug in the host installation/upgrade wizard that occurred when retrying individual hosts. The error message "Configurator not found" might appear and prevent installation or upgrade of all hosts.

Cloudera Bug: OPSAPS-44694

Issues Fixed in Cloudera Manager 5.14.4

Cloudera Manager read-only user can access sensitive cluster information

Fixed an issue where, due to a security vulnerability, a Cloudera Manager read-only user can access sensitive cluster information.

For the latest update on this issue, see the corresponding Knowledge article:

TSB 2018-306: Cloudera Manager Information Disclosure

Cloudera Bug: OPSAPS-45688

Gateway roles and Isilon replication jobs

Fixed a bug where Cloudera Manager did not allow Gateway roles for Isilon clusters, which prevented replications jobs from running.

Cloudera Bug: OPSAPS-46154

Save Changes remains enabled after clicking (possible autocomplete issue)

Google Chrome version 66, released on 28 April 2018, introduced behavior changes which can trigger unintended modification of existing configuration values in Cloudera Manager. If a page contains a password field, the browser aggressively applies auto-complete form data to the first text field and first password form fields. These changes are applied without user input or explicit notification (fields are highlighted).

Cloudera Bug: OPSAPS-46118

Cloudera Backup and Disaster Recovery cannot import Hive metadata

Fixed a bug where Cloudera Backup and Disaster Recovery could not import Hive metadata that included a Partitioned View with a null StorageDescriptor.

Cloudera Bug: OPSAPS-45900

Do not expose configuration history to non-admin users

Users were previously able to go directly to the configuration history and rollback page using the URL, even though the link itself is not visible. This is now protected.

Cloudera Bug: OPSAPS-45781

Cloudera Manager 5 Release Notes

Cloudera Manager should not omit LDAP GROUP_MAPPING passwords from NodeManager

Some sensitive authentication configurations were not emitted for the Yarn Node Manager when using LDAP group mapping. This is now corrected for clusters running CDH 5.5 and higher where encryption is available for those sensitive configurations.

Cloudera Bug: OPSAPS-45440

Errors thrown for Hive replications managed by a single Cloudera Manager

Fixed an issue where an exception is thrown when collecting replication diagnostics for a Hive replication schedule between two clusters managed by a single Cloudera Manager.

Cloudera Bug: OPSAPS-45197

Agent exception "Could not evaluate resource {u'hard_limit': 500, u'soft_limit': 300}: 'unicode' object has no attribute 'items'"

Fixes an issue in Cloudera Manager agent where cgroup soft and hard memory limits were not applied on the host.

Cloudera Bug: OPSAPS-45229

Yarn Application Masters fail when a concurrent Deploy Client Configuration job is executed

Yarn Application Masters fail because the `container-executor.cfg` file is missing when a YARN application is executed concurrently with a Deploy Client Configurations job. The container-executor hierarchy has been moved to `/var/lib/yarn-ce`, which is not modified by the Deploy Client Configuration command.

No changes are required for Cloudera Manager/CDH 5.16.0 and higher. Use the following **Workaround** for Cloudera Manager/CDH 5.15.1 and 5.14.4:

1. On the Cloudera Manager server host, add the following line to the `/etc/default/cloudera-scm-server` file:

```
export CMF_FF_YARN_SAFE_CONTAINER_EXECUTOR_DIR=true
```

2. Restart the Cloudera Manager Server:

```
service cloudera-scm-server restart
```

3. Restart the NodeManager(s) so they can pick up the new configuration change:
 - a. In the Cloudera Manager Admin console, go to the YARN service.
 - b. Click the **Instances** tab.
 - c. Select all hosts with the NodeManager **Role Type**.
 - d. Click **Actions for Selected > Restart**.

Cloudera Bug: OPSAPS-24398

Host installation or upgrade using key based authentication where private key is protected by passphrase does not work

In Cloudera Manager, host installation and upgrade does not work when using key-based host authentication with SSH where the private key is protected by a passphrase.

Cloudera Bug: OPSAPS-45571

Cloudera Manager does not list any CDH package options other than "custom"

When doing a package installation or upgrade of CDH, no versions of CDH5 are shown in the Cloudera Manager Admin Console.

Cloudera Bug: OPSAPS-45732

Issues Fixed in Cloudera Manager 5.14.3

Cross-site scripting vulnerability in Cloudera Manager

Fixed the cross-site vulnerability described in [TSB-2018-285 \(XSS Scripting Vulnerability in Cloudera Manager\)](#). For more information, see [the Known Issue description](#)

Cloudera Bug: DOCS-3186

Backup and Disaster Recovery replication fails between two clusters when data is copied from one encryption zone to another when using CDH 5.12.0 or higher

Fixed an issue where replication fails when two clusters use HDFS HA, and the nameservice names are the same.

Both the source and destination clusters must upgrade to a release that contains the fix.



Important: Cloudera recommends using unique names for the nameservices.

Cloudera Bug: OPSAPS-43556

Fixed in: Cloudera Manager 5.15.0, 5.14.3, 5.13.3

HDFS and Hive Replications may fail after upgrading to CM 5.14.2 on clusters where HDFS HA is enabled

Fixed an issue that caused HDFS and Hive replications to fail if the destination cluster is managed by Cloudera Manager 5.14.2.

For more information, see [TSB-305](#).

Cloudera Bug: OPSAPS-45600

Issues Fixed in Cloudera Manager 5.14.2

Cross-site scripting vulnerability in Cloudera Manager

Fixed the cross-site vulnerability described in [TSB-2018-285 \(XSS Scripting Vulnerability in Cloudera Manager\)](#). For more information, see [the Known Issue description](#)

Cloudera Bug: DOCS-3186

CGroup Hard Memory resource reports incorrect units of measurement

The **Host > Resources** page now reports the correct units for the CGroup Hard Memory resource.

Cloudera Bug: OPSAPS-41473

Scheduled snapshot and replication jobs are skipped

Fixed an issue where scheduled Snapshot and Replication jobs can be skipped if there are too many commands scheduled to run at the same time (more than 100), or if Cloudera Manager has paused due to garbage collection.

Fixed Versions: Cloudera Manager 5.15.0, 5.14.2, 5.13.2

Cloudera Bug: OPSAPS-44153

The Replication performance CSV file header does not match content

Fixed an issue where the header in the HDFS Replication performance file (in CSV format) does not match its contents.

Fixed in: Cloudera Manager 5.15.0, 5.14.2, 5.13.2

Cloudera Bug: OPSAPS-43997

Cloudera Manager 5 Release Notes

BDR Cloud (s3) Replication job fail with delegation token creation issue

Fixed an issue where HDFS to Amazon S3 or Hive to Amazon S3 replication jobs fail when an expired delegation token is used to run a MapReduce/YARN job.

Cloudera Bug: OPSAPS-44008

Cloudera Manager Upgrade (5.14.0 to 5.14.1) stuck at refreshing package metadata - configurator not found

Fixed a bug in the host installation/upgrade wizard that occurred when retrying individual hosts. The error message "Configurator not found" might appear and prevent installation or upgrade of all hosts.

Cloudera Bug: OPSAPS-44694

Hive Replication export step fails with StaleObjectStateException

Fixed an issue where replication commands sometimes fail because of concurrent executions in the database.

Cloudera Bug: OPSAPS-43964

Add "sentry" group in "hadoop.proxyuser.hive.groups" on upgrade to CDH 5.13.0 and higher

For CDH 5.13.0 and higher, when Hive is used with Sentry, Hive Proxy Groups should contain either a wildcard (*) for the Sentry Group for Sentry to work with Hive. If this is not present, Cloudera Manager adds the Sentry Group to Hive Proxy Groups when upgrading CDH.

Cloudera Bug: OPSAPS-43809

Spark CSD should emit proper environment variables for older CDH version 5.4.x

Fixed configuration of Kerberos credentials for the Spark History Server in CDH 5.4.

Cloudera Bug: OPSAPS-43985

Cloudera Manager Server reports 'Too many open files' error

Fixed an issue where collection of diagnostic bundles would lead to slow file descriptor leaks in Cloudera Manager Server that eventually cause the server to become non-operational or unresponsive.

Cloudera Bug: OPSAPS-44922

Table statistics are lost when replicating an existing table without partitions

Cloudera Backup and Disaster Recovery added a workaround for the Hive bug HIVE-15653 and now preserves table statistics for unpartitioned tables during subsequent runs.

Fixed Versions: Cloudera Manager 5.15.0, 5.14.2, 5.13.2

Cloudera Bug: OPSAPS-44070

Issues Fixed in Cloudera Manager 5.14.1

Hive Replications Can Fail Intermittently

This issue, which was also reported in [TSB-2018-276](#), is fixed in CDH 5.13.2. For further details, see the original [Known Issue](#).

Cloudera Manager upgrade workflow incorrectly requires deploying some optional management roles

Fixed an issue that occurred when a user upgrades to Cloudera Manager 5.14.0, and after all the agents have been updated to 5.14.0. The user sees an Upgrade Wizard, and is asked to select where to place the Navigator Metadata Server, Navigator Audit Server, and the Reports Manager roles. The **Continue** button is disabled until user adds one of these roles.

See TSB-289.

Cloudera Bug: OPSAPS-44629

Cloudera Bug: [OPSAPS-44629](#)

Logging issue slows down Hive and HDFS Replication jobs

Fixed an issue that occurred when using Cloudera Manager for cross-cluster replication of HDFS and Hive. Unnecessary warning logs are printed during the initial phase of replication ('copy listing'). These logs were printed for each and every file that Cloudera Backup and Disaster Recovery (BDR) copies. This causes very verbose log output which slows down this phase of the replication job and can affect overall replication times considerably. This happens only when both source and target clusters are running Cloudera Manager 5.14.0.

See TSB-290 for more information.

Fixed Versions: Cloudera Manager 5.15.0, 5.14.1

Cloudera Bug: OPSAPS-44160

Issues Fixed in Cloudera Manager 5.14.0

Transient 502 proxy error displays red popup in Hue

As of CDH 5.12 and later, the Hue Load balancer started throwing transient 502 proxy error that appear as a red popup on Hue web pages. This happens frequently on a secure cluster. The Hue server's `runcpserver.log` file does not log these 502 errors. The error is logged in `/var/log/hue-httpd/error_log`. Cloudera Manager now properly configures the Hue load balancer to avoid these issues.

Cloudera Bug: OPSAPS-43317

Cloudera Manager User Administrator role Error

Fixed a bug that occurred when a Cloudera Manager user with the **User Administrator** goes to the **Settings** page. The **Redaction Parameters for Diagnostic Bundles** parameter displays the following error message:

Invalid JSON configuration: SyntaxError: Unexpected token * in JSON at position 0.

Cloudera Bug: OPSAPS-41286

Cluster upgrade fails in wizard with "There is already a pending command on this entity."

Previously, if the customer selects the **Skip Host Inspector** option, the wizard proceeds to the next page but the command still runs in the background and the user sees the message "There is already a pending command on this entity." Cloudera Manager now ensures that the Host Inspector command is aborted when user selects **Skip Host Inspector**. It is still possible that other commands (scheduled or not) will run at the same time as the upgrade command. In this case, user will still see the same error message.

Cloudera Bug: OPSAPS-43740

Configuration "diff" throws TypeMismatchException

The feature that compares the configuration between multiple services or multiple clusters on the configuration page threw an error in Cloudera Manager 5.13. This is now fixed.

Cloudera Bug: OPSAPS-43130

Send Test Alert response dialog does not work

Since Cloudera Manager 5.12, the ability to send a Test Alert has been broken. This is now fixed.

Cloudera Bug: OPSAPS-42826

Incorrect warnings when configuring HBase Thrift Server in Kerberized cluster

Previously, when configuring the HBase Thrift Server security in a Kerberized cluster, incorrect warnings such as "When HBase is configured to use Kerberos, HBase Thrift Authentication is ignored when Enable TLS/SSL for HBase Thrift Server over HTTP is enabled" would appear. With this fix, these warnings are eliminated.

Cloudera Bug: OPSAPS-41441

Cloudera Manager 5 Release Notes

Historical Disk Usage By User download CSV is not adding the metric type customer is choosing

When creating the "Historical Disk Usage By User" report from the **Cluster** > **Cluster_name** > **Reports** page, the download button did not honor the chosen metric. This is now fixed.

Cloudera Bug: OPSAPS-43474

WebServerMetricCollector breaks

Fixed a bug causing broken Impala monitoring functionality when password authentication is enabled for the Impala Daemon Web server.

Cloudera Bug: OPSAPS-43216

Cloudera Manager LDAP authentication fails

When the External Authentication type is set to LDAP, and the LDAP server is Active Directory, and the LDAP Group Search Base field is configured to be the root of the directory tree, authentication may fail with a `PartialResultException` message in the logs. This exception is now handled so that authentication succeeds.

Cloudera Bug: OPSAPS-43262

Logging in to Cloudera Manager 5.12.x with Internet Explorer 11 fails

Fixed an issue where Internet Explorer defaults to DocMode 7 due to the removal of the `X-UA-Compatible` header in Cloudera Manager 5.12. When the user uses Compatibility Mode in Internet Explorer 9, 10, 11, or Edge, the browser is incorrectly identified as Internet Explorer 7.0 and as a result Cloudera Manager issues an error message (You are using an unsupported browser) and does not allow the user to log in.

Cloudera Bug: OPSAPS-42870

Hive Replication jobs fail intermittently

Fixed a bug where Hive Replications can fail intermittently on the first step. The failure message says: "The remote command failed with error message: Hive Replication Export Step failed".

See TSB-276.

Fixed in: Cloudera Manager 5.14.0, 5.13.1

Cloudera Bug: OPSAPS-43676

Collect HDFS Replication process files in Hive Replication Diagnostic Bundle

The Hive Replication Diagnostic bundle now collects the HDFS Replication step's process log files.

Cloudera Bug: OPSAPS-43583

Hive Replication fails in unsecured clusters when proxy user is not a superuser on replication source

Fixed a bug that caused Hive Replication jobs to fail in unsecured clusters during the transfer step if the proxy user is not a superuser on the replication source.

Cloudera Bug: OPSAPS-43505

lsOf output in host statistics can be huge due to duplicate per-thread entries

Fixed an issue where the output of the `lsOf` command that is captured in support bundles showed open files per-thread as well as per-process. This could cause the size of the support bundle to increase significantly, and in some cases could cause the host statistics collection phase of support bundle collection to fail. Because the per-thread information is generally redundant, it has been removed from the `lsOf` output.

Cloudera Bug: OPSAPS-42646

Cloudera Manager unable to refresh parcel information

Fixes an issue with refreshing parcels. Information about all available parcels was failing to refresh if one of the repositories contained an invalid manifest file.

Cloudera Bug: OPSAPS-43614

Collection of Impala logs should not fail even when some log files are not present

Fixed an issue of missing Impala and Kudu role logs in the diagnostic bundle when their log directories have broken symlinks.

Cloudera Bug: OPSAPS-41194

Transfer Metadata Files step fails when the home directory of the username is encrypted

This fixes a known issue with Hive and HDFS replication where those replications fail if the `/user/$proxyuser` or `/user/hdfs` directories are encrypted. The bug was introduced in Cloudera Manager 5.12 for Hive Replication and in Cloudera Manager 5.13 for HDFS Replication.

Cloudera Bug: OPSAPS-42439

Duplicate replication job on Cloudera Manager restart

This fixes a known issue where a restart of Cloudera Manager that occurs while HDFS or Hive replication jobs are running can cause a duplicate replication job to be started.

Cloudera Bug: OPSAPS-42557

ACLs and external attributes comparison in HDFS and Hive Replication can be incorrect

Fixes a bug where ACLs and external attributes (Xattrs) may be copied over during HDFS and Hive replication even when they have not changed.

Cloudera Bug: OPSAPS-42818

HDFS Replication does not show the progress marker

This fixes an issue in Cloudera Manager in which progress markers for HDFS Replication were not visible in the Admin Console.

Cloudera Bug: OPSAPS-42805

ACLs and external attributes comparison in HDFS and Hive Replication can be incorrect

Fixes a bug where ACLs and external attributes (Xattrs) may be copied over during HDFS and Hive replication even when they have not changed.

Cloudera Bug: OPSAPS-42818

Cloudera Manager fails to start when gateway roles have configuration overrides

Fixes an issue where descriptor generation fails and the Cloudera Management Service fails to start if any Gateway roles have a configuration override.

Cloudera Bug: OPSAPS-43517

Replication fails with 'another command' error when the previous run is aborted

Fixes a bug introduced in Cloudera Manager 5.13 where a replication command is aborted, the next run of the command fails with the following error message: "Another HDFS replication command is in progress".

Cloudera Bug: OPSAPS-42808

Do not include unhealthy hosts when calculating acceptable estimation rate

Fixed an issue affecting the reliable generation of diagnostic bundles when some hosts are in an unhealthy state.

Cloudera Bug: OPSAPS-42720

Duplicate Hive Logs in diagnostic bundles.

Fixes a bug where some Hive logs would be duplicated in diagnostic bundles.

Cloudera Bug: OPSAPS-41611

"100%" health aggregation thresholds no longer permitted

Cloudera Manager now validates that custom services (CSDs) cannot define an invalid threshold for the role health of a service that affects overall service health. Now the thresholds for these health warnings ("percentGreenForGreen" and "percentYellowGreenForYellow ") must be less than 100%.

Cloudera Bug: OPSAPS-42853

Allow validations for HIVE_METASTORE_MAX_MESSAGE_SIZE to be suppressible

Fixed an issue where the validation warning for "Max Message Size for Hive MetaStore" configuration parameter could not be suppressed. The new behavior allows any validations to be suppressed.

Cloudera Bug: OPSAPS-43320

Preserve export.json and distcp-staging path when a different user runs /user replication

This fixes an issue in which a replication job that replicates the `/user` directory in HDFS using the "delete missing" delete policy causes other concurrent replication jobs to fail.

Cloudera Bug: OPSAPS-43050

Spark History Server ignores hadoop.security.auth_to_local mapping

The Spark History Server now respects the configured Kerberos-to-local user name mappings of the HDFS service.

Cloudera Bug: OPSAPS-42577

Impala Queries State are displayed wrong

The Impala queries page now shows the state of running queries correctly instead of calling them Finished.

Cloudera Bug: OPSAPS-43410

Rolling upgrade can fail when high availability services are stopped

Fixed a bug where rolling upgrades of CDH could fail when certain services Cloudera Manager expects to be able to restart without downtime are already stopped. Now, Cloudera Manager tolerates this and starts those services after upgrade as part of the cluster's rolling restart.

Cloudera Bug: OPSAPS-42515

Pagination on Commands page skips some commands

Fixed a bug with the pagination on the **All Recent Commands** page in which commands were not always listed consistently.

Cloudera Bug: OPSAPS-43052

Get all schedules operation is very slow in Cloudera Manager 5.11 and higher

Before Cloudera Manager 5.14, the `getAllSchedules` API returned a list of failed files for all runs of a schedule. This operation slows down execution of the API and has been changed so that a list of failed files is returned only when the view is specified as `full`.

Cloudera Bug: OPSAPS-42587

Improve reporting of Replication progress

Live monitoring of throughput and estimated completion time for HDFS and Hive Replications running using Cloudera Manager Backup and Disaster Recovery (BDR) are now available in log files.

Cloudera Bug: OPSAPS-42756

Remove the recommendation to restart Cloudera Manager when adding trusted Kerberos realm

Changing the list of trusted Kerberos realms no longer requires a restart of Cloudera Manager.

Cloudera Bug: OPSAPS-43237

Inefficiency when checking staleness and host status

Fixed an inefficiency in staleness checking and checking host status that caused slowness in the Cloudera Manager Admin Console.

Cloudera Bug: OPSAPS-42698

NameNode migration timeout while saving namespace of size more than 7GB

Fixes an issue during HDFS NameNode migration between hosts, where the "Saving Namespace" step may timeout if the fsimage size is large.

Cloudera Bug: OPSAPS-38236

Add Diagnostic Events for Agent Upgrades

Added installation and upgrade events for the Cloudera Manager agent software running on a cluster host to the diagnostic bundle.

Cloudera Bug: OPSAPS-42468

Limit the size of Hive log files collected in BDR diagnostic

Hive-related log files included in diagnostic bundle are now be limited to 100MB (before compression).

Cloudera Bug: OPSAPS-41610

Hive partitions not dropped in Hive replication

When partitions are deleted from a Hive table on the source cluster, Cloudera Backup and Disaster Recovery now correctly deletes the same partitions on Hive replication target. This fixes the bug introduced in Cloudera Manager 5.8.

Cloudera Bug: OPSAPS-42613

MR2Params.DEFAULT_MAPREDUCE_REDACTED_PROPERTIES should include Microsoft Azure credentials

Fixes a bug that caused a possible leak of sensitive cloud account information because redaction of these job configuration properties was not enabled by default for MapReduce jobs.

Cloudera Bug: OPSAPS-43753

Cross-site scripting issues fixed

A number of cross-site scripting issues have been identified through internal auditing and have all been fixed.

Cloudera Bug: OPSAPS-42828

Default end time in cluster utilization APIs is not parsed correctly

Fixed an issue where the cluster utilization report APIs did not work if the end time of the report was not specified with the `to` parameter. End time is meant to be an optional parameter, and now such API calls work without the `to` parameter.

Cloudera Bug: OPSAPS-42820

Cloudera Manager 5 Release Notes

Replication fails if the home directory of the username in the schedule is encrypted

Fixed an issue where BDR replication fails if the user directory of the user specified in the **Run as Peer User** or **Run as username** field in the replication schedule is encrypted.

Fixed in: Cloudera Manager 5.14.0, 5.13.1

Cloudera Bug: OPSAPS-42445

Issues Fixed in Cloudera Manager 5.13.3

Cross-site scripting vulnerability in Cloudera Manager

Fixed the cross-site vulnerability described in [TSB-2018-285 \(XSS Scripting Vulnerability in Cloudera Manager\)](#). For more information, see [the Known Issue description](#)

Cloudera Bug: DOCS-3186

Cloudera Manager Server reports 'Too many open files' error

Fixed an issue where collection of diagnostic bundles would lead to slow file descriptor leaks in Cloudera Manager Server that eventually cause the server to become non-operational or unresponsive.

Cloudera Bug: OPSAPS-44922

Backup and Disaster Recovery replication fails between two clusters when data is copied from one encryption zone to another when using CDH 5.12.0 or higher

Fixed an issue where replication fails when two clusters use HDFS HA, and the nameservice names are the same.

Both the source and destination clusters must upgrade to a release that contains the fix.



Important: Cloudera recommends using unique names for the nameservices.

Cloudera Bug: OPSAPS-43556

Fixed in: Cloudera Manager 5.15.0, 5.14.3, 5.13.3

Issues Fixed in Cloudera Manager 5.13.2

Cross-site scripting vulnerability in Cloudera Manager

Fixed the cross-site vulnerability described in [TSB-2018-285 \(XSS Scripting Vulnerability in Cloudera Manager\)](#). For more information, see [the Known Issue description](#)

Cloudera Bug: DOCS-3186

Hive Replications Can Fail Intermittently

This issue, which was also reported in [TSB-2018-276](#), is fixed in CDH 5.13.2. For further details, see the original [Known Issue](#).

Add "sentry" group in "hadoop.proxyuser.hive.groups" on upgrade to CDH 5.13.0 and higher

For CDH 5.13.0 and higher, when Hive is used with Sentry, Hive Proxy Groups should contain either a wildcard (*) for the Sentry Group for Sentry to work with Hive. If this is not present, Cloudera Manager adds the Sentry Group to Hive Proxy Groups when upgrading CDH.

Cloudera Bug: OPSAPS-43809

BDR Cloud (s3) Replication job fail with delegation token creation issue

Fixed an issue where HDFS to Amazon S3 or Hive to Amazon S3 replication jobs fail when an expired delegation token is used to run a MapReduce/YARN job.

Cloudera Bug: OPSAPS-44008

Cloudera Manager User Administrator role Error

Fixed a bug that occurred when a Cloudera Manager user with the **User Administrator** goes to the **Settings** page. The **Redaction Parameters for Diagnostic Bundles** parameter displays the following error message:

Invalid JSON configuration: SyntaxError: Unexpected token * in JSON at position 0.

Cloudera Bug: OPSAPS-41286

Cluster upgrade fails in wizard with "There is already a pending command on this entity."

Previously, if the customer selects the **Skip Host Inspector** option, the wizard proceeds to the next page but the command still runs in the background and the user sees the message "There is already a pending command on this entity." Cloudera Manager now ensures that the Host Inspector command is aborted when user selects **Skip Host Inspector**. It is still possible that other commands (scheduled or not) will run at the same time as the upgrade command. In this case, user will still see the same error message.

Cloudera Bug: OPSAPS-43740

Collect HDFS Replication process files in Hive Replication Diagnostic Bundle

The Hive Replication Diagnostic bundle now collects the HDFS Replication step's process log files.

Cloudera Bug: OPSAPS-43583

Configuration "diff" throws TypeMismatchException

The feature that compares the configuration between multiple services or multiple clusters on the configuration page threw an error in Cloudera Manager 5.13. This is now fixed.

Cloudera Bug: OPSAPS-43130

Historical Disk Usage By User download CSV is not adding the metric type customer is choosing

When creating the "Historical Disk Usage By User" report from the **Cluster > Cluster_name > Reports** page, the download button did not honor the chosen metric. This is now fixed.

Cloudera Bug: OPSAPS-43474

Hive Replication export step fails with StaleObjectStateException

Fixed an issue where replication commands sometimes fail because of concurrent executions in the database.

Cloudera Bug: OPSAPS-43964

Hive Replication fails in unsecured clusters when proxy user is not a superuser on replication source

Fixed a bug that caused Hive Replication jobs to fail in unsecured clusters during the transfer step if the proxy user is not a superuser on the replication source.

Cloudera Bug: OPSAPS-43505

Hive Replication jobs fail intermittently

Fixed a bug where Hive Replications can fail intermittently on the first step. The failure message says: "The remote command failed with error message: Hive Replication Export Step failed".

See TSB-276.

Fixed in: Cloudera Manager 5.14.0, 5.13.1

Cloudera Manager 5 Release Notes

Cloudera Bug: OPSAPS-43676

Collection of Impala logs should not fail even when some log files are not present

Fixed an issue of missing Impala and Kudu role logs in the diagnostic bundle when their log directories have broken symlinks.

Cloudera Bug: OPSAPS-41194

Scheduled snapshot and replication jobs are skipped

Fixed an issue where scheduled Snapshot and Replication jobs can be skipped if there are too many commands scheduled to run at the same time (more than 100), or if Cloudera Manager has paused due to garbage collection.

Fixed Versions: Cloudera Manager 5.15.0, 5.14.2, 5.13.2

Cloudera Bug: OPSAPS-44153

Table statistics are lost when replicating an existing table without partitions

Cloudera Backup and Disaster Recovery added a workaround for the Hive bug HIVE-15653 and now preserves table statistics for unpartitioned tables during subsequent runs.

Fixed Versions: Cloudera Manager 5.15.0, 5.14.2, 5.13.2

Cloudera Bug: OPSAPS-44070

The Replication performance CSV file header does not match content

Fixed an issue where the header in the HDFS Replication performance file (in CSV format) does not match its contents.

Fixed in: Cloudera Manager 5.15.0, 5.14.2, 5.13.2

Cloudera Bug: OPSAPS-43997

WebServerMetricCollector breaks

Fixed a bug causing broken Impala monitoring functionality when password authentication is enabled for the Impala Daemon Web server.

Cloudera Bug: OPSAPS-43216

Spark CSD should emit proper environment variables for older CDH version 5.4.x

Fixed configuration of Kerberos credentials for the Spark History Server in CDH 5.4.

Cloudera Bug: OPSAPS-43985

lsOf output in host statistics can be huge due to duplicate per-thread entries

Fixed an issue where the output of the `lsOf` command that is captured in support bundles showed open files per-thread as well as per-process. This could cause the size of the support bundle to increase significantly, and in some cases could cause the host statistics collection phase of support bundle collection to fail. Because the per-thread information is generally redundant, it has been removed from the `lsOf` output.

Cloudera Bug: OPSAPS-42646

Issues Fixed in Cloudera Manager 5.13.1

Replication fails if the home directory of the username in the schedule is encrypted

Fixed an issue where BDR replication fails if the user directory of the user specified in the **Run as Peer User** or **Run as username** field in the replication schedule is encrypted.

Fixed in: Cloudera Manager 5.14.0, 5.13.1

Cloudera Bug: OPSAPS-42445

Add Diagnostic Events for Agent Upgrades

Added installation and upgrade events for the Cloudera Manager agent software running on a cluster host to the diagnostic bundle.

Cloudera Bug: OPSAPS-42468

Cross-site scripting issues fixed

A number of cross-site scripting issues have been identified through internal auditing and have all been fixed.

Cloudera Bug: OPSAPS-42828

Logging in to Cloudera Manager 5.12.x with Internet Explorer 11 fails

Fixed an issue where Internet Explorer defaults to DocMode 7 due to the removal of the `X-UA-Compatible` header in Cloudera Manager 5.12. When the user uses Compatibility Mode in Internet Explorer 9, 10, 11, or Edge, the browser is incorrectly identified as Internet Explorer 7.0 and as a result Cloudera Manager issues an error message (You are using an unsupported browser) and does not allow the user to log in.

Cloudera Bug: OPSAPS-42870

ACLs and external attributes comparison in HDFS and Hive Replication can be incorrect

Fixes a bug where ACLs and external attributes (Xattrs) may be copied over during HDFS and Hive replication even when they have not changed.

Cloudera Bug: OPSAPS-42818

Transfer Metadata Files step fails when the home directory of the username is encrypted

This fixes a known issue with Hive and HDFS replication where those replications fail if the `/user/$proxyuser` or `/user/hdfs` directories are encrypted. The bug was introduced in Cloudera Manager 5.12 for Hive Replication and in Cloudera Manager 5.13 for HDFS Replication.

Cloudera Bug: OPSAPS-42439

Replication fails with 'another command' error when the previous run is aborted

Fixes a bug introduced in Cloudera Manager 5.13 where a replication command is aborted, the next run of the command fails with the following error message: "Another HDFS replication command is in progress".

Cloudera Bug: OPSAPS-42808

New Advanced Configuration Snippet for overriding the hdfs-site/core-site.xml values for Impala configurations

You can now use the following Advanced Configuration Snippets for `core-site.xml` for all Impalad and Catalog Server roles:

- **Impala Catalog Server Advanced Configuration Snippet (Safety Valve) for core-site.xml**
- **Impala Daemon Advanced Configuration Snippet (Safety Valve) for core-site.xml**

Cloudera Bug: OPSAPS-9586

Improve reporting of Replication progress

Live monitoring of throughput and estimated completion time for HDFS and Hive Replications running using Cloudera Manager Backup and Disaster Recovery (BDR) are now available in log files.

Cloudera Bug: OPSAPS-42756

Cloudera Manager fails to start when gateway roles have configuration overrides

Fixes an issue where descriptor generation fails and the Cloudera Management Service fails to start if any Gateway roles have a configuration override.

Cloudera Bug: OPSAPS-43517

Cloudera Manager 5 Release Notes

Limit the size of Hive log files collected in BDR diagnostic

Hive-related log files included in diagnostic bundle are now be limited to 100MB (before compression).

Cloudera Bug: OPSAPS-41610

Do not include unhealthy hosts when calculating acceptable estimation rate

Fixed an issue affecting the reliable generation of diagnostic bundles when some hosts are in an unhealthy state.

Cloudera Bug: OPSAPS-42720

Duplicate Hive Logs in diagnostic bundles.

Fixes a bug where some Hive logs would be duplicated in diagnostic bundles.

Cloudera Bug: OPSAPS-41611

Hive partitions not dropped in replication target cluster

Fixes a bug introduced in Backup and Disaster Recovery for Cloudera Manager 5.8 where HDFS and Hive replication were not dropping partitions on the target cluster that had been dropped on the source cluster.

Cloudera Bug: OPSAPS-43282

Incorrect warnings when configuring HBase Thrift Server in Kerberized cluster

Previously, when configuring the HBase Thrift Server security in a Kerberized cluster, incorrect warnings such as "When HBase is configured to use Kerberos, HBase Thrift Authentication is ignored when Enable TLS/SSL for HBase Thrift Server over HTTP is enabled" would appear. With this fix, these warnings are eliminated.

Cloudera Bug: OPSAPS-41441

"100%" health aggregation thresholds no longer permitted

Cloudera Manager now validates that custom services (CSDs) cannot define an invalid threshold for the role health of a service that affects overall service health. Now the thresholds for these health warnings ("percentGreenForGreen" and "percentYellowGreenForYellow ") must be less than 100%.

Cloudera Bug: OPSAPS-42853

Duplicate replication job on Cloudera Manager restart

This fixes a known issue where a restart of Cloudera Manager that occurs while HDFS or Hive replication jobs are running can cause a duplicate replication job to be started.

Cloudera Bug: OPSAPS-42557

Allow validations for HIVE_METASTORE_MAX_MESSAGE_SIZE to be suppressible

Fixed an issue where the validation warning for "Max Message Size for Hive MetaStore" configuration parameter could not be suppressed. The new behavior allows any validations to be suppressed.

Cloudera Bug: OPSAPS-43320

Preserve export.json and distcp-staging path when a different user runs /user replication

This fixes an issue in which a replication job that replicates the `/user` directory in HDFS using the "delete missing" delete policy causes other concurrent replication jobs to fail.

Cloudera Bug: OPSAPS-43050

HDFS Replication does not show the progress marker

This fixes an issue in Cloudera Manager in which progress markers for HDFS Replication were not visible in the Admin Console.

Cloudera Bug: OPSAPS-42805

YARN postponed cgroups cleanup could remove hadoop-yarn directory when NodeManager is running

This fixes a bug in which, under certain conditions after a YARN or cluster restart, Cloudera Manager would delete the local directory used by the YARN NodeManager for creation of YARN container cgroups. On systems configured to use cgroups for YARN CPU isolation, this would prevent new YARN workloads from running on specific hosts until the directory was recreated manually.

Cloudera Bug: OPSAPS-40755

Inefficiency when checking staleness and host status

Fixed an inefficiency in staleness checking and checking host status that caused slowness in the Cloudera Manager Admin Console.

Cloudera Bug: OPSAPS-42698

Improved dynamic chunk operation of BDR

For HDFS and Hive replication, copying data in chunks sized by bytes, not a fixed number of files is now the default behavior.

Cloudera Bug: OPSAPS-42682

Remove the recommendation to restart Cloudera Manager when adding trusted Kerberos realm

Changing the list of trusted Kerberos realms no longer requires a restart of Cloudera Manager.

Cloudera Bug: OPSAPS-43237

BDR performance improvement

This fix delivers a big performance boost by parallelizing the process of checking for changes in file permissions, ACLs, and external attributes. Directory permissions are now preserved in the CopyMapper phase instead of the CopyCommitter phase.

Cloudera Bug: OPSAPS-42834

Spark History Server ignores hadoop.security.auth_to_local mapping

The Spark History Server now respects the configured Kerberos-to-local user name mappings of the HDFS service.

Cloudera Bug: OPSAPS-42577

NameNode migration timeout while saving namespace of size more than 7GB

Fixes an issue during HDFS NameNode migration between hosts, where the "Saving Namespace" step may timeout if the fsimage size is large.

Cloudera Bug: OPSAPS-38236

Issues Fixed in Cloudera Manager 5.13.0

CM should report Spark service health as bad when the SparkHistoryServer is down

The Spark service health now reflects the health of its configured roles.

Cloudera Bug: OPSAPS-34698

Timeout between BDR peers increased

The communications timeout between two peer Cloudera Manager servers for communication during BDR replication has been increased to 10 seconds.

Cloudera Bug: OPSAPS-39225

New Impala metrics for hedged reads, JVM heap usage and connection setup queue size

The new metrics are for JVM Heap usage of the Catalog Server and Hedged reads.

Cloudera Bug: OPSAPS-39831, OPSAPS-40212, OPSAPS-40771

ClassCastException during upgrade with custom services (CSD) that support rolling restart

Fixed a bug that caused a `ClassCastException` during upgrade of clusters with custom (CSD) services that explicitly define rolling restart steps for master roles.

Cloudera Bug: OPSAPS-40182

Only one existing NodeManager role marked stale in Cloudera Manager when a new NodeManager is added

Fixes an issue with staleness detection logic. The new staleness detection logic, introduced in 5.11.0, caused a regression where it does not mark all the services stale. This affects rolling restart of the services.

Cloudera Bug: OPSAPS-40332

Proper termination and cleanup of YARN applications during decommissioning

Fixed an issue where YARN applications may not exit gracefully or clean-up properly if hosts are decommissioned while application containers are still running. This could lead to unexpected container/application failure and/or orphaned container processes on YARN worker nodes. Now, when decommissioning a NodeManager or host running a NodeManager, Cloudera Manager first waits for the configured YARN graceful decommission timeout (if applicable), and allows additional time for ResourceManager to properly terminate the NodeManager and its child processes.

Cloudera Bug: OPSAPS-40731

Health Event Startup Policy incorrectly mutes health events

Fixed an issue where all health events are ignored and no alerts get sent when the **Health Event Startup Policy** is set to none.

Cloudera Bug: OPSAPS-40851

Non-ASCII characters in logs caused failure of log collection for diagnostic bundles

Fixed an issue where log files did not get collected correctly as part of diagnostic bundles because they contained non-ASCII characters.

Cloudera Bug: OPSAPS-40949

Exception thrown when searching large log files

Fixed an issue where an exception might be thrown when searching a log when the log file size is larger than 2 GB.

Cloudera Bug: OPSAPS-41710

Exception thrown when accessing a service status page while stopping one of its roles

Fixed an issue where a `NullPointerException` occurred when accessing the service status page in the Cloudera Manager Admin Console while stopping a role.

Cloudera Bug: OPSAPS-41222

Increase Maximum Document Buffer size default to 2047 in Reports Manager

The default upper limit for the Reports Manager's **Maximum Document Buffer Size** configuration parameter has been increased to 2047 MB.

Cloudera Bug: OPSAPS-41358

Spilled memory not populated in Impala queries and in the Spilled Memory tab in the Cluster Utilization report

Fixed an issue where the **Memory Spilled** metric for each Impala query was not correctly populated. This prevented various features that depend on this metric from functioning correctly. For example, the **Spilled Memory** tab of the **Cluster Utilization Report** was not populated with data.

Cloudera Bug: OPSAPS-41447

HSM KMS trust store password is visible in logs

In previous releases, when using the Cloudera Navigator HSM KMS (for use with a Thales or Luna Hardware Security Module), the KMS's trust store password was displayed in some process logs on the KMS hosts. While the trust store password is not considered sensitive, some customers set their trust store password to the same value as their key store password, which is sensitive. In this release, the HSM KMS's trust store password is no longer displayed.

Cloudera Bug: OPSAPS-41626

Hue Load Balancer caused login failure

Fixed an issue where using Apache httpd 2.4 as the Hue load balancer caused Hue logins to fail when TLS is enabled for Hue. This change requires Cloudera Manager to set the **use_x_forwarded_host** parameter in the `hue.ini` file to "true". After upgrading Cloudera Manager, Cloudera Manager will report that the Hue Server and Load Balancer role configurations are stale. If your cluster is affected by this bug, you must restart the Hue service.

Cloudera Bug: OPSAPS-41850

Upgrade is blocked when HMS validation fails

While upgrading to Cloudera Manager 5.12 and CDH 5.12, Cloudera Manager added a new Hive Metastore validation step that is executed as a part of the CDH upgrade command to detect any corruption in the Hive Metastore. In 5.12, when corruption is detected and validation fails, the upgrade command fails and leaves the cluster in an intermediate state, meaning that some services may be started while others are stopped, which is not desirable.

This issue has been fixed in Cloudera Manager 5.12.1. Because it is possible for Hive to tolerate some forms of minor schema corruption, and because it is impossible to run the schema validation tool prior to CDH 5.12, the Hive Metastore validation is now executed as the last step of the upgrade command, which occurs after all the services in the cluster have fully been upgraded and restarted. Although this step still appears to be part of upgrade step, your upgrade is actually complete and you can deal with schema integrity issues following the upgrade. If you see this step failing, exit the upgrade wizard and contact Cloudera Support for assistance repairing any schema corruption. You may then re-run the validation tool manually from the Hive service actions menu after repair is complete to verify a successful repair.

Cloudera Bug: OPSAPS-41936

Default values for `yarn.scheduler.fair.assignmultiple` can block YARN applications from running on CDH 5.5 or CDH 5.6

In some previous versions of Cloudera Manager, the default value of `yarn.scheduler.fair.assignmultiple` was incorrect and can cause problems in CDH 5.5 and 5.6. This has been fixed. Note: if you use CDH versions 5.5.x or 5.6.x and upgrade to CM 5.13, you may see a staleness indicator for the YARN service. The staleness shows `yarn.scheduler.fair.max.assign` as being added. This is expected, and you should restart the YARN service to avoid a serious bug in the Resource Manager ([YARN-4477](#)).

Cloudera Bug: OPSAPS-42118

Add Host Wizard does not apply host template

Fixed an issue where the **Apply Host Template** feature in the **Add Hosts Wizard** did not apply the host template.

Cloudera Bug: OPSAPS-42134

Advanced configuration properties for Hue Load Balancer have no effect

Fixed an issue where Advanced Configuration Snippets for the Hue Load Balancer were not able to override the settings **SSLProtocol** and **SSLCipherSuite**.

Cloudera Bug: OPSAPS-42187

Cloudera Manager 5 Release Notes

Issues with Kafka Active Controller when managing multiple Kafka services

Fixed an issue where Cloudera Manager did not track the Active Controller in Kafka correctly if there were multiple Kafka services running under the same Cloudera Manager instance.

Cloudera Bug: OPSAPS-42216

JDK 7 performance regression for Cloudera Manager-managed deployments of HBase

This change sets the JVM property `ReservedCodeCacheSize` to 256MB in the default JVM startup options for HBase roles. This change attempts to prevent performance issues seen when HBase uses Java 7. The value set is the same as the default when using Java 8.

Cloudera Bug: OPSAPS-42259

NullPointerException on the Command Details page while deploying client configuration

Fixed a `NullPointerException` on the Command Details page that caused the **Deploy Client Configuration** status to not display.

Cloudera Bug: OPSAPS-42390

BDR jobs replicating specific directories can cause other jobs to fail when using certain Delete Policy modes

Fixed an issue where a replication schedule copying the `/user/`, `/user/hdfs` or `/user/$bdr_user` directory with the **Delete Policy** set to delete missing files can cause other replication jobs that run simultaneously to fail.

Cloudera Bug: OPSAPS-42492

Rolling upgrade can fail when high availability services are stopped

Fixed a bug where rolling upgrades of CDH could fail when certain services Cloudera Manager expects to be able to restart without downtime are already stopped. Now, Cloudera Manager tolerates this and starts those services after upgrade as part of the cluster's rolling restart.

Cloudera Bug: OPSAPS-42515

Issues Fixed in Cloudera Manager 5.12.2

Do not include unhealthy hosts when calculating acceptable estimation rate

Fixed an issue affecting the reliable generation of diagnostic bundles when some hosts are in an unhealthy state.

Cloudera Bug: OPSAPS-42720

BDR jobs replicating specific directories can cause other jobs to fail when using certain Delete Policy modes

Fixed an issue where a replication schedule copying the `/user/`, `/user/hdfs` or `/user/$bdr_user` directory with the **Delete Policy** set to delete missing files can cause other replication jobs that run simultaneously to fail.

Cloudera Bug: OPSAPS-42492

Rolling upgrade can fail when high availability services are stopped

Fixed a bug where rolling upgrades of CDH could fail when certain services Cloudera Manager expects to be able to restart without downtime are already stopped. Now, Cloudera Manager tolerates this and starts those services after upgrade as part of the cluster's rolling restart.

Cloudera Bug: OPSAPS-42515

Spark History Server ignores `hadoop.security.auth_to_local` mapping

The Spark History Server now respects the configured Kerberos-to-local user name mappings of the HDFS service.

Cloudera Bug: OPSAPS-42577

Default values for `yarn.scheduler.fair.assignmultiple` can block YARN applications from running on CDH 5.5 or CDH 5.6

In some previous versions of Cloudera Manager, the default value of `yarn.scheduler.fair.assignmultiple` was incorrect and can cause problems in CDH 5.5 and 5.6. This has been fixed. Note: if you use CDH versions 5.5.x or 5.6.x and upgrade to CM 5.13, you may see a staleness indicator for the YARN service. The staleness shows `yarn.scheduler.fair.max.assign` as being added. This is expected, and you should restart the YARN service to avoid a serious bug in the Resource Manager ([YARN-4477](#)).

Cloudera Bug: OPSAPS-42118

[Add Host Wizard does not apply host template](#)

Fixed an issue where the **Apply Host Template** feature in the **Add Hosts Wizard** did not apply the host template.

Cloudera Bug: OPSAPS-42134

[Incorrect warnings when configuring HBase Thrift Server in Kerberized cluster](#)

Previously, when configuring the HBase Thrift Server security in a Kerberized cluster, incorrect warnings such as "When HBase is configured to use Kerberos, HBase Thrift Authentication is ignored when Enable TLS/SSL for HBase Thrift Server over HTTP is enabled" would appear. With this fix, these warnings are eliminated.

Cloudera Bug: OPSAPS-41441

[CM - log refresh takes me to the top of the page every 5 seconds](#)

Fixed an issue in the Install Details dialog that scrolling position was not respected and always moves to the very top. Now it stays in the location that user has scrolled to.

Cloudera Bug: OPSAPS-40836

[YARN postponed cgroups cleanup could remove hadoop-yarn directory when NodeManager is running](#)

This fixes a bug in which, under certain conditions after a YARN or cluster restart, Cloudera Manager would delete the local directory used by the YARN NodeManager for creation of YARN container cgroups. On systems configured to use cgroups for YARN CPU isolation, this would prevent new YARN workloads from running on specific hosts until the directory was recreated manually.

Cloudera Bug: OPSAPS-40755

[New Advanced Configuration Snippet for overriding the hdfs-site/core-site.xml values for Impala configurations](#)

You can now use the following Advanced Configuration Snippets for `core-site.xml` for all Impalad and Catalog Server roles:

- **Impala Catalog Server Advanced Configuration Snippet (Safety Valve) for core-site.xml**
- **Impala Daemon Advanced Configuration Snippet (Safety Valve) for core-site.xml**

Cloudera Bug: OPSAPS-9586

[Allow validations for HIVE_METASTORE_MAX_MESSAGE_SIZE to be suppressible](#)

Fixed an issue where the validation warning for "Max Message Size for Hive MetaStore" configuration parameter could not be suppressed. The new behavior allows any validations to be suppressed.

Cloudera Bug: OPSAPS-43320

[Preserve export.json and distcp-staging path when a different user runs /user replication](#)

This fixes an issue in which a replication job that replicates the `/user` directory in HDFS using the "delete missing" delete policy causes other concurrent replication jobs to fail.

Cloudera Bug: OPSAPS-43050

Cloudera Manager 5 Release Notes

[sentry] Warn when policy files are in use

New Behavior: CM will show a suppressible warning when Sentry Policy files are used when the service can be configured to work with the Sentry Service in that CDH version. Note: HBase Solr Indexer does not currently support the Sentry Service in any CDH version and only supports policy files. Therefore, this validation warning will not be available for HBase Solr Indexer.

Cloudera Bug: OPSAPS-42865

Logging in to Cloudera Manager 5.12.x with Internet Explorer 11 fails

Fixed an issue where Internet Explorer defaults to DocMode 7 due to the removal of the `X-UA-Compatible` header in Cloudera Manager 5.12. When the user uses Compatibility Mode in Internet Explorer 9, 10, 11, or Edge, the browser is incorrectly identified as Internet Explorer 7.0 and as a result Cloudera Manager issues an error message (You are using an unsupported browser) and does not allow the user to log in.

Cloudera Bug: OPSAPS-42870

BDR performance improvement

This fix delivers a big performance boost by parallelizing the process of checking for changes in file permissions, ACLs, and external attributes. Directory permissions are now preserved in the CopyMapper phase instead of the CopyCommitter phase.

Cloudera Bug: OPSAPS-42834

ACLs and external attributes comparison in HDFS and Hive Replication can be incorrect

Fixes a bug where ACLs and external attributes (Xattrs) may be copied over during HDFS and Hive replication even when they have not changed.

Cloudera Bug: OPSAPS-42818

NullPointerException on the Command Details page while deploying client configuration

Fixed a NullPointerException on the Command Details page that caused the **Deploy Client Configuration** status to not display.

Cloudera Bug: OPSAPS-42390

Issues Fixed in Cloudera Manager 5.12.1

Hue Load Balancer caused login failure

Fixed an issue where using Apache httpd 2.4 as the Hue load balancer caused Hue logins to fail when TLS is enabled for Hue. This change requires Cloudera Manager to set the `use_x_forwarded_host` parameter in the `hue.ini` file to "true". After upgrading Cloudera Manager, Cloudera Manager will report that the Hue Server and Load Balancer role configurations are stale. If your cluster is affected by this bug, you must restart the Hue service.

Cloudera Bug: OPSAPS-41850

HSM KMS trust store password is visible in logs

In previous releases, when using the Cloudera Navigator HSM KMS (for use with a Thales or Luna Hardware Security Module), the KMS's trust store password was displayed in some process logs on the KMS hosts. While the trust store password is not considered sensitive, some customers set their trust store password to the same value as their key store password, which is sensitive. In this release, the HSM KMS's trust store password is no longer displayed.

Cloudera Bug: OPSAPS-41626

Maximum Diagnostic Bundle Size reported incorrectly

- The **Send Diagnostic Data** command was erroneously reporting the maximum allowed bundle size when the diagnostic bundle is created using the **By Date Range** option. This message was fixed to show the right allowed bundle size.

- The **Send Diagnostic Data** command displayed unclear error messages when the estimation step fails. This has been fixed to show a clear error message.

Cloudera Bug: OPSAPS-41020

Upgrade is blocked when HMS validation fails

While upgrading to Cloudera Manager 5.12 and CDH 5.12, Cloudera Manager added a new Hive Metastore validation step that is executed as a part of the CDH upgrade command to detect any corruption in the Hive Metastore. In 5.12, when corruption is detected and validation fails, the upgrade command fails and leaves the cluster in an intermediate state, meaning that some services may be started while others are stopped, which is not desirable.

This issue has been fixed in Cloudera Manager 5.12.1. Because it is possible for Hive to tolerate some forms of minor schema corruption, and because it is impossible to run the schema validation tool prior to CDH 5.12, the Hive Metastore validation is now executed as the last step of the upgrade command, which occurs after all the services in the cluster have fully been upgraded and restarted. Although this step still appears to be part of upgrade step, your upgrade is actually complete and you can deal with schema integrity issues following the upgrade. If you see this step failing, exit the upgrade wizard and contact Cloudera Support for assistance repairing any schema corruption. You may then re-run the validation tool manually from the Hive service actions menu after repair is complete to verify a successful repair.

Cloudera Bug: OPSAPS-41936

Spilled memory not populated in Impala queries and in the Spilled Memory tab in the Cluster Utilization report

Fixed an issue where the **Memory Spilled** metric for each Impala query was not correctly populated. This prevented various features that depend on this metric from functioning correctly. For example, the **Spilled Memory** tab of the **Cluster Utilization Report** was not populated with data.

Cloudera Bug: OPSAPS-41447

Increase Maximum Document Buffer size default to 2047 in Reports Manager

The default upper limit for the Reports Manager's **Maximum Document Buffer Size** configuration parameter has been increased to 2047 MB.

Cloudera Bug: OPSAPS-41358

Exception thrown when searching large log files

Fixed an issue where an exception might be thrown when searching a log when the log file size is larger than 2 GB.

Cloudera Bug: OPSAPS-41710

Proper termination and cleanup of YARN applications during decommissioning

Fixed an issue where YARN applications may not exit gracefully or clean-up properly if hosts are decommissioned while application containers are still running. This could lead to unexpected container/application failure and/or orphaned container processes on YARN worker nodes. Now, when decommissioning a NodeManager or host running a NodeManager, Cloudera Manager first waits for the configured YARN graceful decommission timeout (if applicable), and allows additional time for ResourceManager to properly terminate the NodeManager and its child processes.

Cloudera Bug: OPSAPS-40731

Viewing agent logs in Internet Explorer renders incorrectly

When viewing agent logs in Internet Explorer, the logs may render as HTML instead of text. This has been fixed by forcing browsers to render the log page as text.

Cloudera Bug: OPSAPS-41122

Threads CPU time and Total Time are not populated in Impala Query Monitoring

The **Thread CPU Time** and **Total Time** attributes of Impala Query Monitoring are now correctly populated.

Cloudera Bug: OPSAPS-39647

Issues Fixed in Cloudera Manager 5.12.0

Renaming an empty directory on Amazon S3 when using DynamoDB with S3Guard may fail

Fixed an issue where renaming an empty directory on Amazon S3 when using DynamoDB with S3Guard may fail

Cloudera bug: CDH-51869

Apache bug: HADOOP-14036

Changes to the `/tmp` directory causes Hue failures

Changing permissions or removing files in the `/tmp` directory can cause Hue to fail because the Hue Kerberos ticket cache file, located by default in the `/tmp` directory, is no longer available. The default location has been changed to `/var/run/hue`.

You can change the location of the Hue Kerberos ticket cache by changing the **Hue Kerberos Credentials Cache Directory** property. In Cloudera Manager, go to the Hue service, select the **Configuration** tab, and search for the property. Restart the Hue service to enable the change.

After upgrading Cloudera Manager you must also restart the Hue service.

Cloudera Bug: OPSAPS-26101

Exception when setting `maintenanceOwners` to null

Fixed how null values for the `maintenanceOwners` parameter are handled when creating clusters with the Cloudera Manager API.

Cloudera Bug: OPSAPS-32880

Upgrade continues when there are hosts with warnings

Fixed an issue where an upgrade continues even if hosts have a warning. Host warnings must be resolved before an upgrade can complete.

Cloudera Bug: OPSAPS-40945

Address renewal requirements for delegation tokens in DR

Fixed an issue where BDR replication jobs fail on a Kerberos enabled cluster if the job duration is longer than the renewal interval for the hdfs delegation token. With this fix, both the delegation token and Kerberos ticket are renewed until the max lifetime of token/ticket (default value is 7 days). This enables longer replications without needing to bring down the source cluster.

Cloudera Bug: OPSAPS-32811

`hostname` parameter is not passed to Impala catalog role

[IMPALA-5253](#) contained a security fix for clusters using Impala with TLS (SSL) security enabled. This fix was also made in several maintenance versions of CDH that require you to upgrade Cloudera Manager. If you upgrade to a CDH version with this fix without upgrading Cloudera Manager, Impala will not function when TLS is enabled for Impala. You should upgrade Cloudera Manager first if you want to move to a CDH version with the security fix.

This issue affects upgrades of Cloudera Manager and CDH to version 5.11.1.

There are two ways you can workaround this issue:

- Upgrade to one of the following versions of Cloudera Manager before upgrading CDH:
 - 5.13.0
 - 5.12.1
 - 5.11.2
 - 5.10.2
 - 5.9.3

– 5.8.5

-or-

- *Before upgrading CDH*, set the `-hostname` option to the fully-qualified domain name of the Catalog Server using the **Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)** configuration property:

```
-hostname=fully-qualified-domain-name of Impala Catalog Server
```

(To set this property, in Cloudera Manager, go to the Impala service, select the **Configuration** tab and search for the property.

CDH versions with the Impala security fix:

- 5.11.1
- 5.10.2
- 5.9.3
- 5.8.5

Cloudera bug: OPSAPS-41218

Hive Replication fails with OTHER_ERROR with duplicate view/table names

Fixed an issue where Hive replication fails if the same name is used for a table on the target cluster which is also used for a view on the source cluster. With this change, Hive replication properly handles all cases where a view and table on source/target clusters have the same name.

Cloudera Bug: OPSAPS-36302

Underscores in LDAP domain names not allowed

Fixed an issue where LDAP domain names could not contain underscores. This fix affects Cloudera Manager, Hue and Cloudera Navigator.

Cloudera Bug: OPSAPS-40487

Only use healthy HDFS/Hive hosts for launching replication jobs

BDR Replication Host Selection Policy has been updated. The process that launches and coordinates a HDFS/Hive replication job will now only run on the following hosts:

- Hosts that run any role of the HDFS/Hive Service (for HDFS or Hive replication)
- Hosts that have a Non-Gateway role
- Hosts where the health status is in the GOOD or CONCERNING state with preference given to GOOD
- Hosts that are whitelisted, if configured

Cloudera Bug: OPSAPS-40040

Low watermark value for Memstore Flush default is incorrect

For CDH versions 5.8 and higher, the **Low Watermark for Memstore Flush** configuration parameter is associated with the HBase parameter `hbase.regionserver.global.memstore.lowerLimit`.

This value represents the fullness threshold of the memstore as a percentage of memstore capacity. The default value for this parameter was incorrectly set too low at `.38`. This can cause severe under utilization of the memstore.

The default has been corrected to be `.95`. When upgrading to a version of Cloudera Manager with this fix, if the value was previously set to the old default of `.38`, it will automatically be increased to the new default, which may cause Cloudera Manager to mark your HBase service as having a stale configuration, requiring a restart.

Additionally, if an existing **Low Watermark for Memstore Flush** configuration parameter has a value $\leq .9$, it will be flagged as a configuration warning.

Cloudera Manager 5 Release Notes

Fixed in: Cloudera Manager 5.12, 5.11.1, 5.10.2, 5.9.3, 5.8.5

Cloudera Bug: OPSAPS-38468

Accessing Sqoop2 with Hue fails

Fixed an issue where accessing Sqoop2 with Hue fails with the following error: Sqoop error: Could not get connectors.

CDH 5.5.0 or higher clusters running a Hue Service require a restart of Hue services because of a new configuration file. This new configuration file is required for Hue's Sqoop2 Application functions in secure settings, but it will show up whether or not Sqoop or security settings are in use.

Cloudera Bug: OPSAPS-27286

Make Oozie Load-balancer URL mandatory when enabling Oozie HA

Fixed an issue when the **Oozie Load Balancer** field in the Oozie configuration is not set and Oozie high availability is in use (two or more Oozie Server roles are deployed). This now creates a validation error and Cloudera Manager does not let you start or restart Oozie until the problem is fixed.

Fixed in: Cloudera Manager 5.12, 5.11.2, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-39372

Wire encryption configuration properties

HDFS client data-transfer encryption cipher related properties are now published to HTTPFs. Previously they were emitted just to HDFS Gateways, which caused HTTPFs performance to suffer when Data Transfer Encryption was turned on. When upgrading to Cloudera Manager 5.12, all HTTPFs instances will be marked stale, and a restart of those roles is required for these properties to take effect.

Cloudera Bug: OPSAPS-39981

Missing Sentry configuration files

Fixed an issue where some configuration files (core-site.xml) were missing from Sentry configuration.

Cloudera Bug: OPSAPS-33710

LoggingOutInterceptor log level

To enhance security, when API debugging is enabled, the entire API request and response is no longer printed at the DEBUG log level. To see the request/response, set the logging level to TRACE. Note that this may expose sensitive information in the logs.

Cloudera Bug: OPSAPS-34538

Failed health checks because of deprecated ntpdc command

Fixed an issue where the `ntpdc` command was used in the host Clock Offset health test even if the command is deprecated for an operating system.

Cloudera Bug: OPSAPS-38268

Reports Manager throws ArrayIndexOutOfBoundsException when indexing fsimage

Fixes an issue where an `ArrayIndexOutOfBoundsException` can be thrown by Reports Manager.

Fixed in Cloudera Manager 5.12, 5.11.2, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40914

Spark History Server password handling

Improved password handling for the Spark History Server. After upgrading to Cloudera Manager 5.12, Cloudera Manager may consider the Spark service as having a stale configuration due to this change, requiring a restart.

Cloudera Bug: OPSAPS-40298

Agent process directories re-mounted on agent restart

Fixed an issue where restarting an agent on newer Linux distributions such as Ubuntu 16.04 or Red Hat 7 would result in process run directories being remounted and lost.

Cloudera Bug: OPSAPS-41259

Process secrets no longer protected with file permissions

Fixed an issue where secret data, including passwords and other data, can be exposed in the `/var/run/cloudera-scm-agent/process-name/proc.json` or `/var/run/cloudera-scm-agent/process-name/config.zip` files because these files are world readable. See TSB-235 for more information.

Fixed in: Cloudera Manager 5.12, 5.11, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40536

Configuration files can be downloaded from the UI

Certain files that may contain sensitive information can no longer be downloaded from the Processes tab in Cloudera Manager. If access to those files is required, you must log into the appropriate host.

Cloudera Bug: OPSAPS-40309

Command download filename uniqueness

Fixed an issue where downloading the command output did not generate a unique file name.

Cloudera Bug: OPSAPS-40249

StopRunner is not honouring the keyAdministrator Role Activities

Fixed an issue where a CSD descriptor was not respecting the correct authorization for start/stop when specifying a service-level stopRunner (which enables graceful shutdown). This could cause some user roles to be able to start but not stop some services.

Cloudera Bug: OPSAPS-41030

Hue Load Balancer SSL Handshake error

Fixed an issue with the SSL handshake in Hue when Apache httpd 2.4 or higher is used as the load balancer. The Hue load balancer previously set the `ProxyPreserveHost` directive to `On`, when it should have been set to `Off`. This causes problems making SSL connections when using Apache httpd 2.4 or higher. The error caused problems when verifying the CN, which older versions of Apache httpd did not encounter because they did not properly verify the CN.

When upgrading Cloudera Manager, the Hue load balancer may be marked as having a stale configuration. If you are experiencing issues connecting to Hue with SSL, restart the Hue service to update the configuration.

Cloudera Bug: OPSAPS-40700

LDAP trust failures not showing in stack traces

When an external user authenticating via LDAP or Active Directory fails to login to Cloudera Manager, the failure stack trace is now logged in the Cloudera Manager server log to help debug the failure.

Cloudera Bug: OPSAPS-40065

getRoles filtering is not working with Cloudera Manager API version 11 and higher

Fixed a bug that broke backwards compatibility with Cloudera Manager API version 11 (introduced with the Cloudera Manager 5.6 release) for the following endpoint:

```
/api//clusters/clusterName/services/serviceName/roles?filter=type==ROLE_TYPE
```

Cloudera Bug: OPSAPS-40184

Cluster templates don't include role level configuration

Fixed an issue with Cluster Templates where role level configuration was not getting exported. This was causing cluster to fail when trying user try to import a cluster template exported from a cluster on which HDFS HA was enabled.

Cloudera Bug: OPSAPS-31528

Dynamic Resource Pool drop-down options are not visible when using Internet Explorer 11

Fixed a typo on the **Dynamic Resource Pools** page that prevents Internet Explorer from rendering the add/edit resource pool dialog box.

Cloudera Bug: OPSAPS-40522

Cloudera Manager Server not validating agent hostnames with TLS

Fixed an issue where, when using TLS, hostnames were not verified during agent to server communication. To fix this, Cloudera Manager has a new setting: **Verify Agent Hostname Against Certificate**. When enabled, agent hostnames must match the hostname in the agent certificate, which is validated when agents send heartbeats to the Cloudera Manager Server.

Cloudera Bug: OPSAPS-40125

User Administrator role cannot configure S3Guard

Fixed an issue where the user administrator role was disallowed from managing S3 configuration.

Cloudera Bug: OPSAPS-40083

Allow diagnostic bundle estimation timeout to be configurable

Sometimes on a large or busy cluster, collection of diagnostic data using the **By Date Range** option can fail due to a timeout during estimation step of the diagnostic bundle collection. You can now configure both the host level estimation timeout and the overall estimation timeout using Java options. These options can be set on the Cloudera Manager Server host I in the `/etc/default/cloudera-scm-server` file by adding the following to the line that begins with `export CMF_JAVA_OPTS:`

- `-Dcom.cloudera.RoleLogEstimator.maxEstimateTimeoutSeconds=number of seconds` (If not specified, the default is 90 seconds)
- `-Dcom.cloudera.RoleLogEstimator.estimateTimeoutPerHostSeconds=number of seconds` (If not specified, default is 60 seconds)



Note:

The value of `-Dcom.cloudera.RoleLogEstimator.maxEstimateTimeoutSeconds=number of seconds` must be `>=` the value of `-Dcom.cloudera.RoleLogEstimator.estimateTimeoutPerHostSeconds=number of seconds`.

Restart the Cloudera Manager server for the updated Java flags to take effect:

```
sudo service cloudera-scm-server restart
```

Fixed in: Cloudera Manager 5.12, 5.11.2, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40987

Peak Memory Usage report for Impala does not display data

Fixed an issue where the **Impala Peak Memory Usage** page in the Cluster Utilization Report did not display any data in Cloudera Manager.

Fixed in: Cloudera Manager 5.12, 5.11.2, 5.9.3

Cloudera Bug: OPSAPS-39836

NullPointerException when accessing the configuration page for Alerts

Fixed an issue where a NullPointerException occurred when viewing the configuration page for alerts.

Fixed in Cloudera Manager 5.12, 5.11.2, 5.10.2

Cloudera Bug: OPSAPS-39296

Issues Fixed in Cloudera Manager 5.11.2

Agent process directories re-mounted on agent restart

Fixed an issue where restarting an agent on newer Linux distributions such as Ubuntu 16.04 or Red Hat 7 would result in process run directories being remounted and lost.

Cloudera Bug: OPSAPS-41259

hostname parameter is not passed to Impala catalog role

[IMPALA-5253](#) contained a security fix for clusters using Impala with TLS (SSL) security enabled. This fix was also made in several maintenance versions of CDH that require you to upgrade Cloudera Manager. If you upgrade to a CDH version with this fix without upgrading Cloudera Manager, Impala will not function when TLS is enabled for Impala. You should upgrade Cloudera Manager first if you want to move to a CDH version with the security fix.

This issue affects upgrades of Cloudera Manager and CDH to version 5.11.1.

There are two ways you can workaround this issue:

- Upgrade to one of the following versions of Cloudera Manager before upgrading CDH:
 - 5.13.0
 - 5.12.1
 - 5.11.2
 - 5.10.2
 - 5.9.3
 - 5.8.5

-or-

- *Before upgrading CDH*, set the `-hostname` option to the fully-qualified domain name of the Catalog Server using the **Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)** configuration property:

```
-hostname=fully-qualified-domain-name of Impala Catalog Server
```

(To set this property, in Cloudera Manager, go to the Impala service, select the **Configuration** tab and search for the property.

CDH versions with the Impala security fix:

- 5.11.1
- 5.10.2
- 5.9.3
- 5.8.5

Cloudera bug: OPSAPS-41218

Maximum Diagnostic Bundle Size reported incorrectly

- The **Send Diagnostic Data** command was erroneously reporting the maximum allowed bundle size when the diagnostic bundle is created using the **By Date Range** option. This message was fixed to show the right allowed bundle size.
- The **Send Diagnostic Data** command displayed unclear error messages when the estimation step fails. This has been fixed to show a clear error message.

Cloudera Bug: OPSAPS-41020

Allow diagnostic bundle estimation timeout to be configurable

Sometimes on a large or busy cluster, collection of diagnostic data using the **By Date Range** option can fail due to a timeout during estimation step of the diagnostic bundle collection. You can now configure both the host level estimation timeout and the overall estimation timeout using Java options. These options can be set on the Cloudera Manager Server host I in the `/etc/default/cloudera-scm-server` file by adding the following to the line that begins with `export CMF_JAVA_OPTS:`

- `-Dcom.cloudera.RoleLogEstimator.maxEstimateTimeoutSeconds=number of seconds` (If not specified, the default is 90 seconds)
- `-Dcom.cloudera.RoleLogEstimator.estimateTimeoutPerHostSeconds=number of seconds` (If not specified, default is 60 seconds)



Note:

The value of `-Dcom.cloudera.RoleLogEstimator.maxEstimateTimeoutSeconds=number of seconds` must be \geq the value of `-Dcom.cloudera.RoleLogEstimator.estimateTimeoutPerHostSeconds=number of seconds`.

Restart the Cloudera Manager server for the updated Java flags to take effect:

```
sudo service cloudera-scm-server restart
```

Fixed in: Cloudera Manager 5.12, 5.11.2, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40987

Reports Manager throws `ArrayIndexOutOfBoundsException` when indexing fsimage

Fixes an issue where an `ArrayIndexOutOfBoundsException` can be thrown by Reports Manager.

Fixed in Cloudera Manager 5.12, 5.11.2, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40914

Proper termination and cleanup of YARN applications during decommissioning

Fixed an issue where YARN applications may not exit gracefully or clean-up properly if hosts are decommissioned while application containers are still running. This could lead to unexpected container/application failure and/or orphaned container processes on YARN worker nodes. Now, when decommissioning a NodeManager or host running a NodeManager, Cloudera Manager first waits for the configured YARN graceful decommission timeout (if applicable), and allows additional time for ResourceManager to properly terminate the NodeManager and its child processes.

Cloudera Bug: OPSAPS-40731

Peak Memory Usage report for Impala does not display data

Fixed an issue where the **Impala Peak Memory Usage** page in the Cluster Utilization Report did not display any data in Cloudera Manager.

Fixed in: Cloudera Manager 5.12, 5.11.2, 5.9.3

Cloudera Bug: OPSAPS-39836

Make Oozie Load-balancer URL mandatory when enabling Oozie HA

Fixed an issue when the **Oozie Load Balancer** field in the Oozie configuration is not set and Oozie high availability is in use (two or more Oozie Server roles are deployed). This now creates a validation error and Cloudera Manager does not let you start or restart Oozie until the problem is fixed.

Fixed in: Cloudera Manager 5.12, 5.11.2, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-39372

NullPointerException when accessing the configuration page for Alerts

Fixed an issue where a NullPointerException occurred when viewing the configuration page for alerts.

Fixed in Cloudera Manager 5.12, 5.11.2, 5.10.2

Cloudera Bug: OPSAPS-39296

Issues Fixed in Cloudera Manager 5.11.1

All required fonts are now installed by Cloudera Manager

Fixed an issue where Cloudera Manager made requests to googleapi.com to download some of its required fonts, which fails if the browser does not have Internet access. Cloudera Manager now includes all of the necessary fonts.

Cloudera Bug: OPSAPS-40609.

Automated Cloudera Manager installer fails on Ubuntu 16.04

Fixed an issue where running the `cloudera-manager-installer.bin` installer file (as described in the [documentation](#)) fails on Ubuntu 16.04 LTS (Xenial).

Cloudera Bug: DOCS-2037.

Unable to connect to Oozie with curl

Fixed an issue where some Linux distributions could not connect to Oozie with `curl` through HTTPS when DHE-based ciphers are present.

Cloudera Bug: OPSAPS-40407

getRoles filtering is not working with Cloudera Manager API version 11 and higher

Fixed a bug that broke backwards compatibility with Cloudera Manager API version 11 (introduced with the Cloudera Manager 5.6 release) for the following endpoint:

```
/api//clusters/clusterName/services/serviceName/roles?filter=type==ROLE_TYPE
```

Cloudera Bug: OPSAPS-40184

Exception when setting maintenanceOwners to null

Fixed how null values for the `maintenanceOwners` parameter are handled when creating clusters with the Cloudera Manager API.

Fixed in: Cloudera Manager 5.12, 5.11.1, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-32880

HBase configuration is supplied for Hive when HBase service is not selected

Fixed an issue where Cloudera Manager provided configuration for the `hive-site.xml` file even if the HBase Service setting in Hive is not selected, which could cause unnecessary errors when Hive-on-Spark attempts to connect to HBase.

Cloudera Manager 5 Release Notes

Cloudera Manager now correctly emits the HBase-related configuration in the `hive-site.xml` only when Hive is dependent on HBase.

This issue has been fixed in the following releases:

- CM5.8.5 and higher maintenance releases,
- CM5.9.2 and higher maintenance releases,
- CM5.10.2 and higher maintenance releases,
- CM5.11.1 and higher maintenance releases

When you upgrade Cloudera Manager from an older version to one of the fixed versions, many services may have stale configurations. You must restart all the stale services in order for them to receive the updated configuration in the `hive-site.xml` file.

Cloudera Bug: OPSAPS-39021

Accessing Sqoop2 with Hue fails

Fixed an issue where accessing Sqoop2 with Hue fails with the following error: Sqoop error: Could not get connectors.

CDH 5.5.0 or higher clusters running a Hue Service require a restart of Hue services because of a new configuration file. This new configuration file is required for Hue's Sqoop2 Application functions in secure settings, but it will show up whether or not Sqoop or security settings are in use.

Cloudera Bug: OPSAPS-27286

Exception when setting maintenanceOwners to null

Fixed how null values for the `maintenanceOwners` parameter are handled when creating clusters with the Cloudera Manager API.

Cloudera Bug: OPSAPS-32880

Underscores in LDAP domain names not allowed

Fixed an issue where LDAP domain names could not contain underscores. This fix affects Cloudera Manager, Hue and Cloudera Navigator.

Cloudera Bug: OPSAPS-40487

Only use healthy HDFS/Hive hosts for launching replication jobs

BDR Replication Host Selection Policy has been updated. The process that launches and coordinates a HDFS/Hive replication job will now only run on the following hosts:

- Hosts that run any role of the HDFS/Hive Service (for HDFS or Hive replication)
- Hosts that have a Non-Gateway role
- Hosts where the health status is in the GOOD or CONCERNING state with preference given to GOOD
- Hosts that are whitelisted, if configured

Cloudera Bug: OPSAPS-40040

CSD graceful service stop second step failure

Fixed an issue in the stop command execution for CSD services authored to use a service-level graceful shutdown. The stop command could be shown with the second step (forced kill) marked as failed when all roles were already stopped. This issue also affects shutdown of Kafka using Cloudera Manager 5.11.0 and can impact Kafka upgrades.

Fixed in: Cloudera Manager 5.11.1, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40002

Metric missing from Workload Summary for Kudu

Fixed an issue where the `total_kudu_rows_upserted_rate_across_kudu_replicas` metric was not included in the Workload Summary for Kudu.

Cloudera Bug: OPSAPS-40551

Low watermark value for Memstore Flush default is incorrect

For CDH versions 5.8 and higher, the **Low Watermark for Memstore Flush** configuration parameter is associated with the HBase parameter `hbase.regionserver.global.memstore.lowerLimit`.

This value represents the fullness threshold of the memstore as a percentage of memstore capacity. The default value for this parameter was incorrectly set too low at `.38`. This can cause severe under utilization of the memstore.

The default has been corrected to be `.95`. When upgrading to a version of Cloudera Manager with this fix, if the value was previously set to the old default of `.38`, it will automatically be increased to the new default, which may cause Cloudera Manager to mark your HBase service as having a stale configuration, requiring a restart.

Additionally, if an existing **Low Watermark for Memstore Flush** configuration parameter has a value $\leq .9$, it will be flagged as a configuration warning.

Fixed in: Cloudera Manager 5.12, 5.11.1, 5.10.2, 5.9.3, 5.8.5

Cloudera Bug: OPSAPS-38468

Failed health checks because of deprecated ntpdc command

Fixed an issue where the `ntpdc` command was used in the host Clock Offset health test even if the command is deprecated for an operating system.

Cloudera Bug: OPSAPS-38268

Dynamic Resource Pool drop-down options are not visible when using Internet Explorer 11

Fixed a typo on the **Dynamic Resource Pools** page that prevents Internet Explorer from rendering the add/edit resource pool dialog box.

Cloudera Bug: OPSAPS-40522

Cluster templates don't include role level configuration

Fixed an issue with Cluster Templates where role level configuration was not getting exported. This was causing cluster to fail when trying user try to import a cluster template exported from a cluster on which HDFS HA was enabled.

Cloudera Bug: OPSAPS-31528

Graceful shutdown of Kafka brokers does not work as expected

Fixed an issue where the new **Graceful Shutdown Timeout** configuration property does not work as expected. As a result, Kafka takes an additional 30 seconds (by default) to shut down, but will still only have 30 seconds to complete its controlled shutdown, before Cloudera Manager forcibly shuts down Kafka brokers regardless of the configured timeout.

Fixed in Cloudera Manager 5.11.1.

Cloudera Bug: OPSAPS-40106.

Issues Fixed in Cloudera Manager 5.11.0

ResourceManager state store recovery may fail.

Fixed an issue where ResourceManager state store recovery may fail. The YARN service will become stale when a cluster containing a YARN service with a ZooKeeper dependency set is upgraded to Cloudera Manager 5.10.1+ or 5.11+. Restart the YARN service normally or with a rolling restart.

Cloudera Manager 5 Release Notes

Cloudera Bug: [OPSAPS-38908](#)

Cloudera Bug: OPSAPS-38908

Privilege Escalation in Cloudera Manager

Fixed an issue where a read-only Cloudera Manager user can discover the usernames of other users and elevate the privileges of another user. A user cannot elevate their own privilege.

Cloudera Bug: OPSAPS-34124

The `cloudera-server-db fast_stop` fails on Debian-based systems

Fixed an issue where the `cloudera-server-db fast_stop` command fails on non-sytemd Debian systems like Debian 7 with a command not found error.

Cloudera Bug: OPSAPS-39393

Cluster template export exports Autoconfiguration settings

Fixed an issue where the API to export cluster templates exported settings that were auto-configured. Auto-configured settings should not be exported in the cluster template.

Cloudera Bug: OPSAPS-39031

Pauses in Cloudera Manager after adding peer

Fixed an issue that caused pauses and slow performance of the Cloudera Manager Admin Console occur after creating a peer.

Fixed in Cloudera Manager 5.11

Cloudera Bug: OPSAPS-38868.

Max heap is wrong for ZooKeeper

Fixed an issue where the JVM Heap Memory Usage chart for ZooKeeper displayed the wrong value for maximum memory.

Cloudera Bug: OPSAPS-38821

Error starting the NodeManager when the `/tmp` directory is mounted as `noexec`

Fixed an issue where an error occurs when starting the NodeManager because the `/tmp` directory is mounted as `noexec`.

Cloudera Bug: OPSAPS-33053

Cloudera Manager cannot delete principals from AD when regenerating principals

Fixed an issue where generating Kerberos credentials failed when Active Directory properties like OU have a space in them or are too long to fit in a line.

Cloudera Bug: OPSAPS-37100

The Cloudera Manager Python API may experience connection errors

Python 2.7 and later enables certificate chain validation by default, and specifying a custom SSL context with an CA certificate may be required to prevent connection errors. The `ApiResource` constructor in the Cloudera Manager Python API now takes an optional `"ssl_context"` argument where a custom SSL context can be provided. This allow clients to specify the CA certificate of the API endpoint's certificate.

Cloudera Bug: OPSAPS-31699

The service `cloudera-scm-server force_start` command fails with Postgres SQL

Fixed an issue where a force start of the Cloudera Manager server failed if Cloudera Manager used PostgreSQL.

Cloudera Bug: OPSAPS-32915

Parcels cannot be removed from Cloudera Manager

Fixed an issue where parcels could not be removed from the Cloudera Manager database when the `manifest.json` file is not available in the remote repository or removed from the remote location.

Cloudera Bug: OPSAPS-34964

Service Monitor cannot retrieve Navigator Metadata Server metrics when SSL is enabled

Fixed an issue where the Service Monitor cannot retrieve Navigator Metadata Server metrics when SSL is enabled for Navigator Metadata Server.

Fixed in 5.11, 5.10.1, 5.9.2, 5.7.6

Cloudera Bug: OPSAPS-38865

Warning about suspicious non-ASCII punctuation character in KMS ACLs

When setting up the KMS ACLs for HDFS encryption, the warning "Suspicious non-ASCII punctuation character "" at character position 9118" displays. This warning can be safely ignored, and editing any KMS ACL in Cloudera Manager will make the warning disappear. This bug fix removes this warning.

Fixed in Cloudera Manager 5.11.0, 5.10.1

Cloudera Bug: OPSAPS-38670

Kafka fails when configured with Sentry and an old Kafka version

Fixed an issue that occurred when Kafka was configured to use Sentry for authorization, and Kafka was not version 2.1 or higher, which caused Kafka to fail with the following exception: `java.lang.ClassNotFoundException: org.apache.sentry.kafka.authorizer.SentryKafkaAuthorizer`.

Cloudera Bug: OPSAPS-37907

Fixed in: Cloudera Manager 5.11.0, 5.10.1

Invalid regex needs agent hard_restarted

Fixed Cloudera Manager agent problem that requires hard restart when an invalid regular expression is used to specify tables in Hive replication.

Fixed in: Cloudera Manager 5.11.0, 5.10.1, 5.9.2

Cloudera Bug: OPSAPS-38946 and OPSAPS-36974

Issue with views after BDR Hive Replication

Fixed an issue where a change to a Hive view on a source cluster was not reflected in the target cluster after Hive Replication with the Force Overwrite option selected.

Cloudera Bug: OPSAPS-39034

Sentry Database upgrades

During a CDH upgrade, the Sentry Database Upgrade step will be performed only when necessary. Previously, the Sentry Database Upgrade step always ran.

Cloudera Bug: OPSAPS-25405

Hive Replication fails when Impala is SSL enabled but Hadoop services are not

Fixed an issue where a `NullPointerException` is thrown when using Hive Replication with SSL enabled on an Impala service but not on other Hadoop services.

Cloudera Bug: OPSAPS-38700

Fixed in Cloudera Manager 5.11, 5.10.1, 5.9.2

Issues Fixed in Cloudera Manager 5.10.2

Underscores in LDAP domain names not allowed

Fixed an issue where LDAP domain names could not contain underscores. This fix affects Cloudera Manager, Hue and Cloudera Navigator.

Cloudera Bug: OPSAPS-40487

Low watermark value for Memstore Flush default is incorrect

For CDH versions 5.8 and higher, the **Low Watermark for Memstore Flush** configuration parameter is associated with the HBase parameter `hbase.regionserver.global.memstore.lowerLimit`.

This value represents the fullness threshold of the memstore as a percentage of memstore capacity. The default value for this parameter was incorrectly set too low at `.38`. This can cause severe under utilization of the memstore.

The default has been corrected to be `.95`. When upgrading to a version of Cloudera Manager with this fix, if the value was previously set to the old default of `.38`, it will automatically be increased to the new default, which may cause Cloudera Manager to mark your HBase service as having a stale configuration, requiring a restart.

Additionally, if an existing **Low Watermark for Memstore Flush** configuration parameter has a value $\leq .9$, it will be flagged as a configuration warning.

Fixed in: Cloudera Manager 5.12, 5.11.1, 5.10.2, 5.9.3, 5.8.5

Cloudera Bug: OPSAPS-38468

Heap dump files don't contain the correct process ID

Fixed an issue where if Sentry crashed with an OutOfMemory error, the generated dump file name was not unique, and therefore the file could be overwritten on subsequent occurrences of the error. This has been fixed to include the PID in the dump file name, which is the standard practice for other CDH services.

Cloudera Bug: OPSAPS-40600

CSD graceful service stop second step failure

Fixed an issue in the stop command execution for CSD services authored to use a service-level graceful shutdown. The stop command could be shown with the second step (forced kill) marked as failed when all roles were already stopped. This issue also affects shutdown of Kafka using Cloudera Manager 5.11.0 and can impact Kafka upgrades.

Fixed in: Cloudera Manager 5.11.1, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40002

getRoles filtering is not working with Cloudera Manager API version 11 and higher

Fixed a bug that broke backwards compatibility with Cloudera Manager API version 11 (introduced with the Cloudera Manager 5.6 release) for the following endpoint:

```
/api//clusters/clusterName/services/serviceName/roles?filter=type==ROLE_TYPE
```

Cloudera Bug: OPSAPS-40184

Only use healthy HDFS/Hive hosts for launching replication jobs

BDR Replication Host Selection Policy has been updated. The process that launches and coordinates a HDFS/Hive replication job will now only run on the following hosts:

- Hosts that run any role of the HDFS/Hive Service (for HDFS or Hive replication)
- Hosts that have a Non-Gateway role
- Hosts where the health status is in the GOOD or CONCERNING state with preference given to GOOD

- Hosts that are whitelisted, if configured

Cloudera Bug: OPSAPS-40040

Dynamic Resource Pool drop-down options are not visible when using Internet Explorer 11

Fixed a typo on the **Dynamic Resource Pools** page that prevents Internet Explorer from rendering the add/edit resource pool dialog box.

Cloudera Bug: OPSAPS-40522

HBase configuration is supplied for Hive when HBase service is not selected

Fixed an issue where Cloudera Manager provided configuration for the `hive-site.xml` file even if the HBase Service setting in Hive is not selected, which could cause unnecessary errors when Hive-on-Spark attempts to connect to HBase. Cloudera Manager now correctly emits the HBase-related configuration in the `hive-site.xml` only when Hive is dependent on HBase.

This issue has been fixed in the following releases:

- CM5.8.5 and higher maintenance releases,
- CM5.9.2 and higher maintenance releases,
- CM5.10.2 and higher maintenance releases,
- CM5.11.1 and higher maintenance releases

When you upgrade Cloudera Manager from an older version to one of the fixed versions, many services may have stale configurations. You must restart of all the stale services in order for them to receive the updated configuration in the `hive-site.xml` file.

Cloudera Bug: OPSAPS-39021

The `cloudera-server-db fast_stop` fails on Debian-based systems

Fixed an issue where the `cloudera-server-db fast_stop` command fails on non-sytemd Debian systems like Debian 7 with a command not found error.

Cloudera Bug: OPSAPS-39393

Exception when setting `maintenanceOwners` to null

Fixed how null values for the `maintenanceOwners` parameter are handled when creating clusters with the Cloudera Manager API.

Cloudera Bug: OPSAPS-32880

Cannot upload large diagnostic bundles

Fixed an issue where large diagnostic bundles, such as ones greater than 4 GB, fail to upload.

Cloudera Bug: OPSAPS-40100

NullPointerException when accessing the configuration page for Alerts

Fixed an issue where a NullPointerException occurred when viewing the configuration page for alerts.

Fixed in Cloudera Manager 5.12, 5.11.2, 5.10.2

Cloudera Bug: OPSAPS-39296

Issues Fixed in Cloudera Manager 5.10.1

ResourceManager state store recovery may fail.

Fixed an issue where ResourceManager state store recovery may fail. The YARN service will become stale when a cluster containing a YARN service with a ZooKeeper dependency set is upgraded to Cloudera Manager 5.10.1+ or 5.11+. Restart the YARN service normally or with a rolling restart.

Cloudera Bug: [OPSAPS-38908](#)

Cloudera Bug: OPSAPS-38908

Privilege Escalation in Cloudera Manager

Fixed an issue where a read-only Cloudera Manager user can discover the usernames of other users and elevate the privileges of another user. A user cannot elevate their own privilege.

Cloudera Bug: OPSAPS-34124

Warning about suspicious non-ASCII punctuation character in KMS ACLs

When setting up the KMS ACLs for HDFS encryption, the warning "Suspicious non-ASCII punctuation character "" at character position 9118" displays. This warning can be safely ignored, and editing any KMS ACL in Cloudera Manager will make the warning disappear. This bug fix removes this warning.

Fixed in Cloudera Manager 5.11.0, 5.10.1

Cloudera Bug: OPSAPS-38670

Hive Replication fails when Impala is SSL enabled but Hadoop services are not

Fixes a bug where some Hadoop services did not recognize the Impala SSL configuration, which caused Hive replications to fail when SSL is enabled. To enable SSL, see: [Configuring TLS/SSL for HDFS, YARN and MapReduce](#).

This support also ensures that connection to Impala is successful when SSL is enabled even if Kerberos is not enabled.

Fixed in: Cloudera Manager 5.10.1, 5.9.2, 5.8.4

Cloudera Bug: OPSAPS-38700

Service Monitor cannot retrieve Navigator Metadata Server metrics when SSL is enabled

Fixed an issue where the Service Monitor cannot retrieve Navigator Metadata Server metrics when SSL is enabled for Navigator Metadata Server.

Fixed in 5.11, 5.10.1, 5.9.2, 5.7.6

Cloudera Bug: OPSAPS-38865

Spark CSD may cause issues in yarn-cluster mode

Spark client configuration has been fixed to avoid application failures in certain cases when launching applications in YARN cluster mode.

Fixed in Cloudera Manager 5.10.1, 5.9.2

Cloudera Bug: OPSAPS-39118

Invalid regex needs agent hard_restarted

Fixed Cloudera Manager agent problem that requires hard restart when an invalid regular expression is used to specify tables in Hive replication.

Fixed in: Cloudera Manager 5.11.0, 5.10.1, 5.9.2

Cloudera Bug: OPSAPS-38946 and OPSAPS-36974

Kafka fails when configured with Sentry and an old Kafka version

Fixed an issue that occurred when Kafka was configured to use Sentry for authorization, and Kafka was not version 2.1 or higher, which caused Kafka to fail with the following exception: `java.lang.ClassNotFoundException: org.apache.sentry.kafka.authorizer.SentryKafkaAuthorizer`.

Cloudera Bug: OPSAPS-37907

Fixed in: Cloudera Manager 5.11.0, 5.10.1

Hive Replication fails when Impala is SSL enabled but Hadoop services are not

Fixed an issue where a NullPointerException is thrown when using Hive Replication with SSL enabled on an Impala service but not on other Hadoop services.

Cloudera Bug: OPSAPS-38700

Fixed in Cloudera Manager 5.11, 5.10.1, 5.9.2

Issues Fixed in Cloudera Manager 5.10.0

Importing cluster templates fails

Fixes an issue where the Cluster Import command fails after updating parcel repositories and when the parameter `addRepositories=true` is included in the command and there are no pre-configured repositories in the cluster.

Cloudera Bug: OPSAPS-33059

DB Oracle client libraries cannot be found on host due to missing libaio1

On a parcels install, Hue can be set up with an Oracle Database. This requires placing required Oracle packages in a specific location. See [Hue Custom Databases](#).

Cloudera Bug: OPSAPS-34339

CM - error connecting to an Oracle DB with a Package install for HUE's DB

On a packages install, Hue can be set up with an Oracle Database and proceed past the Test Connection page in the setup wizard. This requires placing required Oracle packages in a specific location. See [Hue Custom Databases](#).

Cloudera Bug: OPSAPS-34608

Support non-public schemas with PostgreSQL

Cloudera Manager used to work with only 'public' schema in PostgreSQL. With this change, Cloudera Manager now supports custom schema names. Cloudera Manager relies on the schema search path in PostgreSQL, which can be set for a user or database. For more information, see <https://www.postgresql.org/docs/current/static/ddl-schemas.html>.

Cloudera Bug: OPSAPS-33752

Truststore password shouldn't be required

Trust store passwords are not required. This fixes a bug where some Cloudera components erroneously reported that a password was required.

Cloudera Bug: OPSAPS-38183

Server error deleting a host - Cannot delete or update a parent row: a foreign key constraint fails

Fixes a database constraint violation that occurs when you use Cloudera Manager to delete a host that has cluster-wide configurations. (Kerberos configurations, for example.)

Cloudera Bug: OPSAPS-38150

Hide unlicensed configurations instead of disabling them

Cloudera Express no longer shows configurations that requires a license to manage.

Cloudera Bug: OPSAPS-38214

Various improvements in the Pools page

Many small usability issues have been fixed in the resource pools page:

- You can now click on an input row and go directly to that input field in the edit dialog.
- The columns where you enter minimum and maximum values for Virtual cores and Memory now appear under headings **Min Resources** and **Max Resources**

- On the **Create Resource Pool** page, the **Configuration Sets** tab has been renamed to **Resource Limits**.

Cloudera Bug: OPSAPS-38133

Override the SMON client configuration so the MapReduce job does not produce compressed output

This fixes an issue where YARN utilization report may not work when MapReduce job output is specified as "Compressed" and a compression codec is used for it.

Cloudera Bug: OPSAPS-38460

Excess load on Cloudera Manager database

Fixed an issue where the load on the Cloudera Manager database increases significantly on clusters with 100-200 nodes or more and when there are host-level configuration overrides. This has been resolved by using a smaller default value for the `default_batch_fetch_size` property in the hibernate configuration.

Fixed in: Cloudera Manager 5.10.0, 5.9.2, 5.8.4, 5.7.6

Cloudera Bug: OPSAPS-37985

Dollar sign '\$' character in password results in an IllegalArgumentException

Fixed an issue where CSD-based services could not use certain special characters like '\$' without getting an exception. This is most frequently an issue for passwords.

Cloudera Bug: OPSAPS-37542

Incorrect Scheduling Policy of Resource Pool Configuration in YARN

If a user did not make any resource policy change on an existing cluster, the default scheduling policy used is FAIR, but the Cloudera Manager Admin console incorrectly displays the policy as **DRF**. This is now fixed to show **FAIR**. When the user creates new clusters, the default scheduling policy will be **DRF**.

Cloudera Bug: OPSAPS-37478

Cannot select time period in custom charts in C5.9

The quick time selection (30m, 1h, etc) on custom dashboards now work correctly.

Cloudera Bug: OPSAPS-37190

Increase the default value of Safemode Minimum DataNodes property to 1

In Cloudera Manager 5.9 and lower, the HDFS property **Safemode Minimum DataNodes** has a default value of 0. This property configures the number of DataNodes that must register with the NameNode before the NameNode exits safe mode. This property now has a default value of 1 in CDH 5.9 and higher

In an empty cluster (that is, a cluster with no HDFS data) where the property is set to 0, this has the side effect of extending the duration of startup safe mode by about 30 seconds. For non-empty clusters, this change has no visible effect. This change only applies to new CDH 5.9 and higher clusters created with Cloudera Manager 5.10 or higher. Existing CDH 5.9 clusters are not affected; pre-CDH 5.9 clusters that are upgraded to CDH 5.9 or higher are also unaffected.

Cloudera Bug: OPSAPS-36988

Turn off HSTS header in Hue Load Balancer

Only the Hue server now only generates the HSTS HTTP header. Previously, both the Hue server and Hue Load Balancer generated the HSTS HTTP header.

Cloudera Bug: OPSAPS-35536

AWS ID being used for UUID, which can cause duplicate hosts

Prevent duplicate hosts from being reported in Cloudera Manager due to use of the AWS instance ID as a unique host identifier. Cloudera Manager now uses a random number as the host unique identifier instead. This prevents a duplicate host issue when host is running on a cloud service such as AWS and the host is added using the Cloudera Manager wizard and manually added later.

Cloudera Bug: OPSAPS-35282

Cloudera Manager corrupts passwords and other attributes

Previously, some users found that '*****' were inserted in the database literally. Cloudera Manager now performs a check on this special string prevents it from being stored in the Cloudera Manager database.

Cloudera Bug: OPSAPS-34252

Database wizard error message when the database does not exist is misleading

Cloudera Manager now generates a proper error message when the database does not exist and the username and password do exist. Cloudera Manager now displays the message "Able to find the Database server, but not the specified database. Please check if the database name is correct and make sure that the user can access the database."

Cloudera Bug: OPSAPS-33915

Hive Replication: Make parameters replication on by default

Parameters of databases, tables, partitions, and indexes are replicated by default during Hive replications. You can disable this replication. See [Replication of Parameters](#).

Cloudera Bug: OPSAPS-38197

NFS Gateway daemon in HDFS needs SSL server configs

The Web server of the HDFS NFS Gateway that is used to publish metrics and view configurations is now covered by SSL protection when enabled across the HDFS service. This fixes an issue that caused a failure when starting the Web interface for the NFS gateway.

Cloudera Bug: OPSAPS-33202

CM server error - JsonMappingException

Fixed an issue where a race condition can happen when getting the log links for unhealthy roles and then causes a null pointer exception for the Cloudera Manager homepage.

Cloudera Bug: OPSAPS-32804

"Restart needed" button doesn't appear after changing the server's private key password in the KeyTrustee Server service configuration

When passwords change, the stale configuration icons displays. When you click one of these icons, the **Stale Configurations** page displays but the page does not indicate which passwords have changed.

Cloudera Bug: OPSAPS-29146

Issues Fixed in Cloudera Manager 5.9.3

Reports Manager throws ArrayIndexOutOfBoundsException when indexing fsimage

Fixes an issue where an `ArrayIndexOutOfBoundsException` can be thrown by Reports Manager.

Fixed in Cloudera Manager 5.12, 5.11.2, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40914

Make Oozie Load-balancer URL mandatory when enabling Oozie HA

Fixed an issue when the **Oozie Load Balancer** field in the Oozie configuration is not set and Oozie high availability is in use (two or more Oozie Server roles are deployed). This now creates a validation error and Cloudera Manager does not let you start or restart Oozie until the problem is fixed.

Fixed in: Cloudera Manager 5.12, 5.11.2, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-39372

Allow diagnostic bundle estimation timeout to be configurable

Sometimes on a large or busy cluster, collection of diagnostic data using the **By Date Range** option can fail due to a timeout during estimation step of the diagnostic bundle collection. You can now configure both the host level estimation timeout and the overall estimation timeout using Java options. These options can be set on the Cloudera Manager Server host I in the `/etc/default/cloudera-scm-server` file by adding the following to the line that begins with `export CMF_JAVA_OPTS:`

- `-Dcom.cloudera.RoleLogEstimator.maxEstimateTimeoutSeconds=number of seconds` (If not specified, the default is 90 seconds)
- `-Dcom.cloudera.RoleLogEstimator.estimateTimeoutPerHostSeconds=number of seconds` (If not specified, default is 60 seconds)



Note:

The value of `-Dcom.cloudera.RoleLogEstimator.maxEstimateTimeoutSeconds=number of seconds` must be \geq the value of `-Dcom.cloudera.RoleLogEstimator.estimateTimeoutPerHostSeconds=number of seconds`.

Restart the Cloudera Manager server for the updated Java flags to take effect:

```
sudo service cloudera-scm-server restart
```

Fixed in: Cloudera Manager 5.12, 5.11.2, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40987

Low watermark value for Memstore Flush default is incorrect

For CDH versions 5.8 and higher, the **Low Watermark for Memstore Flush** configuration parameter is associated with the HBase parameter `hbase.regionserver.global.memstore.lowerLimit`.

This value represents the fullness threshold of the memstore as a percentage of memstore capacity. The default value for this parameter was incorrectly set too low at `.38`. This can cause severe under utilization of the memstore.

The default has been corrected to be `.95`. When upgrading to a version of Cloudera Manager with this fix, if the value was previously set to the old default of `.38`, it will automatically be increased to the new default, which may cause Cloudera Manager to mark your HBase service as having a stale configuration, requiring a restart.

Additionally, if an existing **Low Watermark for Memstore Flush** configuration parameter has a value $\leq .9$, it will be flagged as a configuration warning.

Fixed in: Cloudera Manager 5.12, 5.11.1, 5.10.2, 5.9.3, 5.8.5

Cloudera Bug: OPSAPS-38468

Exception when setting maintenanceOwners to null

Fixed how null values for the `maintenanceOwners` parameter are handled when creating clusters with the Cloudera Manager API.

Cloudera Bug: OPSAPS-32880

CSD graceful service stop second step failure

Fixed an issue in the stop command execution for CSD services authored to use a service-level graceful shutdown. The stop command could be shown with the second step (forced kill) marked as failed when all roles were already stopped. This issue also affects shutdown of Kafka using Cloudera Manager 5.11.0 and can impact Kafka upgrades.

Fixed in: Cloudera Manager 5.11.1, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40002

Heap dump files don't contain the correct process ID

Fixed an issue where if Sentry crashed with an OutOfMemory error, the generated dump file name was not unique, and therefore the file could be overwritten on subsequent occurrences of the error. This has been fixed to include the PID in the dump file name, which is the standard practice for other CDH services.

Cloudera Bug: OPSAPS-40600

Only use healthy HDFS/Hive hosts for launching replication jobs

BDR Replication Host Selection Policy has been updated. The process that launches and coordinates a HDFS/Hive replication job will now only run on the following hosts:

- Hosts that run any role of the HDFS/Hive Service (for HDFS or Hive replication)
- Hosts that have a Non-Gateway role
- Hosts where the health status is in the GOOD or CONCERNING state with preference given to GOOD
- Hosts that are whitelisted, if configured

Cloudera Bug: OPSAPS-40040

Process secrets no longer protected with file permissions

Fixed an issue where secret data, including passwords and other data, can be exposed in the `/var/run/cloudera-scm-agent/process-name/proc.json` or `/var/run/cloudera-scm-agent/process-name/config.zip` files because these files are world readable. See TSB-235 for more information.

Fixed in: Cloudera Manager 5.12, 5.11, 5.10.2, 5.9.3

Cloudera Bug: OPSAPS-40536

Peak Memory Usage report for Impala does not display data

Fixed an issue where the **Impala Peak Memory Usage** page in the Cluster Utilization Report did not display any data in Cloudera Manager.

Fixed in: Cloudera Manager 5.12, 5.11.2, 5.9.3

Cloudera Bug: OPSAPS-39836

Validator for `hive.metastore.server.max.message.size` can break Hive Metastore startup

The recommended value for **Max Message Size for Hive MetaStore** configuration parameter should be at least 10% of the value for the **Java Heap Size of Hive Metastore Server in Bytes** parameter, but should never exceed 2147483647B. Previously, the validator for **Max Message Size for Hive MetaStore** was showing an incorrect value in its validation message. This validator is now fixed to show the correct recommended value.

Cloudera Bug: OPSAPS-38434

`getRoles` filtering is not working with Cloudera Manager API version 11 and higher

Fixed a bug that broke backwards compatibility with Cloudera Manager API version 11 (introduced with the Cloudera Manager 5.6 release) for the following endpoint:

```
/api//clusters/clusterName/services/serviceName/roles?filter=type==ROLE_TYPE
```

Cloudera Bug: OPSAPS-40184

Issues Fixed in Cloudera Manager 5.9.2

Privilege Escalation in Cloudera Manager

Fixed an issue where a read-only Cloudera Manager user can discover the usernames of other users and elevate the privileges of another user. A user cannot elevate their own privilege.

Cloudera Bug: OPSAPS-34124

Excess load on Cloudera Manager database

Fixed an issue where the load on the Cloudera Manager database increases significantly on clusters with 100-200 nodes or more and when there are host-level configuration overrides. This has been resolved by using a smaller default value for the `default_batch_fetch_size` property in the hibernate configuration.

Fixed in: Cloudera Manager 5.10.0, 5.9.2, 5.8.4, 5.7.6

Cloudera Bug: OPSAPS-37985

Service Monitor cannot retrieve Navigator Metadata Server metrics when SSL is enabled

Fixed an issue where the Service Monitor cannot retrieve Navigator Metadata Server metrics when SSL is enabled for Navigator Metadata Server.

Fixed in 5.11, 5.10.1, 5.9.2, 5.7.6

Cloudera Bug: OPSAPS-38865

Invalid regex needs agent hard_restarted

Fixed Cloudera Manager agent problem that requires hard restart when an invalid regular expression is used to specify tables in Hive replication.

Fixed in: Cloudera Manager 5.11.0, 5.10.1, 5.9.2

Cloudera Bug: OPSAPS-38946 and OPSAPS-36974

Spark CSD may cause issues in yarn-cluster mode

Spark client configuration has been fixed to avoid application failures in certain cases when launching applications in YARN cluster mode.

Fixed in Cloudera Manager 5.10.1, 5.9.2

Cloudera Bug: OPSAPS-39118

Dollar sign '\$' character in password results in an IllegalArgumentException

Fixed an issue where CSD-based services could not use certain special characters like '\$' without getting an exception. This is most frequently an issue for passwords.

Cloudera Bug: OPSAPS-37542

Hive Replication fails when Impala is SSL enabled but Hadoop services are not

Fixed an issue where a NullPointerException is thrown when using Hive Replication with SSL enabled on an Impala service but not on other Hadoop services.

Cloudera Bug: OPSAPS-38700

Fixed in Cloudera Manager 5.11, 5.10.1, 5.9.2

Issues Fixed in Cloudera Manager 5.9.1

Cannot select time period in custom charts in C5.9

The quick time selection (30m, 1h, and so on) on custom dashboards does not work in 5.9.x.

Cloudera Bug: OPSAPS-37190

Hive table Views do not get restored from S3

When creating a Hive Replication schedule that copies Hive data from S3 and you select the **Reference Data From Cloud** option, Hive table Views are not restored correctly and result in a Null Pointer Exception when querying data from the view.

Cloudera Bug: OPSAPS-37549

Cloudera Manager blocks from HS2 enabling both LDAP and Kerberos authorization

HS2 now supports LDAP and Kerberos authentication on the same instance for CDH 5.7.0 or higher. Previously, this was considered an error.

Cloudera Bug: OPSAPS-34310

Support non-public schemas

Cloudera Manager used to work with only the `public` schema in Postgresql. With this change, Cloudera Manager now supports custom schema names. Cloudera Manager relies on the schema search path in postgresql, which can be set for a user, database, etc: See <https://www.postgresql.org/docs/current/static/ddl-schemas.html>

Cloudera Bug: OPSAPS-33752

Alert Publisher throws "unsigned 32bit value" error after long uptime

Fixes an issue where the Alert Publisher throws "unsigned 32bit value" error after long uptime. TimeTick is an extension of UnsignedInteger32 type. TimeTick represents time in 1/100 seconds. Its value range is 0 to 4294967295. Activity Monitor's uptime is a Java long type that can exceed the range of TimeTicks. Before this fix, long uptime could cause the Alert Publisher to throw the "unsigned 32bit value" error.

Cloudera Bug: OPSAPS-35564

Issues Fixed in Cloudera Manager 5.9

Configuration changes marking too many roles as stale

When upgrading to Cloudera Manager 5.9 or later, YARN will become stale due to the following properties changing to be emitted to the Resource Manager role only:

- `yarn.scheduler.fair.user-as-default-queue`
- `yarn.scheduler.fair.preemption`
- `yarn.scheduler.fair.sizebasedweight`
- `yarn.scheduler.fair.assignmultiple`
- `yarn.scheduler.fair.locality.threshold.node`
- `yarn.scheduler.fair.locality.threshold.rack`

This fixes a problem where changes to these configurations would mark too many roles as stale. You can safely ignore staleness to any role caused by the removal of these configuration parameters.

Cloudera Bug: OPSAPS-32610

Impala `load_catalog_in_background` configuration should be set to *false* by default in Cloudera Manager

Changed default value for `load_catalog_in_background` to *false*. This fixes catalog scalability issues observed in Impala since v2.2.

Cloudera Manager 5 Release Notes

Cloudera Bug: OPSAPS-35869

Oozie first run fails with custom principals

Fixed an issue in Cloudera Manager 5.8.1 where the first time Oozie is run it fails if using Kerberos custom principals.

Cloudera Bug: OPSAPS-35822

Cluster export fails when service configuration is invalid

Fixed an issue in the export cluster template code path where it was failing because of stale configuration settings in the Cloudera Manager database. This can occur when configurations are deprecated on older CDH releases.

Cloudera Bug: OPSAPS-35586

Add gateway role to Kafka and move dependencies' client configs into Kafka's client configs

Cloudera recommends that you place your Kafka host in the same logical cluster as your Sentry host. However, if you deploy Kafka as a separate logical cluster, you can deploy a dummy Sentry service on Kafka's logical cluster with an override for `Sentry-site.xml` to point to the Sentry service on first logical cluster, and then it can be turned off. This is a workaround that allows Cloudera Manager to generate the appropriate Sentry client configurations for Kafka.

Cloudera Bug: OPSAPS-35369

Sentry Upgrade command is not retrievable

If the Sentry Upgrade command fails, you can now retry the command.

Cloudera Bug: OPSAPS-35365

Impala Breakpad script does not convert exponentials into decimals and leads to errors

Fixed an issue where the Impala Breakpad script fails if you try to collect more than 10 MB of dumps from a single role.

Cloudera Bug: OPSAPS-34971

HDFS-S3: Incremental backup support

HDFS-S3 replication supports incremental backups using snapshot diffs. HDFS-S3 replication does not support incremental restore using snapshot diff.

Cloudera Bug: OPSAPS-34987

Fix css in navigation

Fixed several small issues:

1. Do not show role link again on the role instances page.
2. Do not bold the decommission state.
3. Do not use long display name for hosts; use short display name.
4. Show the role link on the role health test page.
5. Show *Today* rather than the long time label.
6. Show only month, date when the year is the same.
7. Make the icons in navigation line up.
8. Make the icons in navigation not so wide.
9. Change all line-height settings to multiples of 4.
10. Add back the border below the title.

Cloudera Bug: OPSAPS-35262

Fix sorting on the instances page

Fixed the initial sort order on the instances page.

Cloudera Bug: OPSAPS-35268

Nightly Charts are broken due to gridster refactoring

There is a permission issue on the Cluster Status page where the grid is never enabled, and on the View Page where the grid is always enabled for all users. The former change allows the user to customize the Cluster status page. The latter change ensures that users do not get errors when customizing specific view pages.

Cloudera Bug: OPSAPS-34526

Redact s3 credential properties in MR by default.

AWS S3 credentials, if specified in the job configuration by setting `fs.s3a.access.key` and `fs.s3a.secret.key`, were shown as clear text in MapReduce UIs. The credentials are now redacted by default in all MapReduce UIs.

Cloudera Bug: OPSAPS-36840

Latest Cloudera Manager Java client does not work on older Cloudera Manager

Fixed an issue in Cloudera Manager client versions 5.8, 5.8.1, and 5.8.2 where `ApiCollectDiagnosticDataArguments` was incompatible with Cloudera Manager versions lower than 5.8.

Cloudera Bug: OPSAPS-36315

Comment for command retry in the cm_api Python github needs to be updated

Fixed the API command documentation to include the `canRetry` attribute. Added a new method, `ApiCommand.getCanRetry()` and deprecated the method `ApiCommand.isCanRetry()`.

Cloudera Bug: OPSAPS-36009

cm_api changes for HDFS Cloud replication schedule args

python api changes for hdfs cloud replication.

Cloudera Bug: OPSAPS-35892

Defaults to 1st Impala service in the cluster when there are multiple Impalas

Fixed an issue that prevented the display of Impala usage on the **Cluster Utilization** page when there are two Impala services in the same cluster. The UI shows usage for the first Impala service only.

Cloudera Bug: OPSAPS-34495

Staleness check page popped up 30s ~ 60s after clicking the icon.

Fixed an issue where a stale configuration page would take a lot of time to load for a large cluster.

Cloudera Bug: OPSAPS-34497

Oozie points to older sharelib even after running sharelib install command

Fixed the Install Oozie Share Lib action so that the Oozie service is informed that there is a new shared library installed. This eliminates the need for a separate manual restart.

Cloudera Bug: OPSAPS-26825

Cloudera Manager set catalogd default jvm memory to 4G can cause out of memory error on upgrade to Cloudera Manager 5.7+

After upgrading to 5.7 or later, you might see a reduced Java heap maximum on Impala Catalog Server due to a change in its default value. Upgrading from Cloudera Manager lower than 5.7 to Cloudera Manager 5.8.2 no longer causes any effective change in the Impala Catalog Server Java Heap size.

When upgrading from Cloudera Manager 5.7 or later to Cloudera Manager 5.8.2, if the Impala Catalog Server Java Heap Size is set at the default (4GB), it is automatically changed to either 1/4 of the physical RAM on that host, or 32GB, whichever is lower. This can result in a higher or a lower heap, which could cause additional resource contention or out of memory errors, respectively.

Cloudera Bug: OPSAPS-34039

Agent uses USER env var to run everything if it is set

Cloudera Manager Server now runs services using the user ID with which that single-user mode is configured. If single-user mode is configured on the agents to use the user *cloudera-scm*, then an attempt to run Hive as the user *hive* fails.

Cloudera Bug: OPSAPS-35621

In LZO enabled clusters, the cluster utilization feature fails

YARN usage aggregation job now runs successfully on clusters where LZO compression is being used.

Cloudera Bug: OPSAPS-33356

HDFS Snapshot policy is selecting unhealthy host to run on

When selecting a role to run HDFS Snapshot command, Cloudera Manager selects a non-decommissioned host in Active status. Also, hosts in maintenance mode have a lower priority than the ones in active state.

Cloudera Bug: OPSAPS-33144

Discourage CSDs from producing client config files that pollute the root dir

For CSD development, there are now deprecation warnings in the CLI validation tool and the Maven validation plugin. These indications provide guidance for future behavior changes in the support of CSD services in Cloudera Manager. A deprecation warning is shown when applicable, but does not affect the outcome of the validation.

Cloudera Bug: OPSAPS-33510

Ensure that one-off processes cannot re-execute.

One-off requests, such as "Inspect Hosts," do not re-execute.

Cloudera Bug: OPSAPS-34255

Hive Replication Command should update copy the Serde properties correctly

Hive Replication now replicates the Serde Properties and also copies the corresponding HDFS file.

Cloudera Bug: OPSAPS-34354

Cloudera Manager complains "Unable to parse the input XML" when you enter an working XML for the "Fair Scheduler XML Advanced Configuration Snippet (Safety Valve)"

Passing a legal XML value to Fair Scheduler XML Advanced Configuration Snippet (Safety Valve) no longer causes a parsing error.

Cloudera Bug: OPSAPS-34329

If files excluded by exclusion filters are renamed, they are not replicated

Customers using Incremental HDFS replication had an issue where if an excluded file (through exclusion filters) is renamed to an included file, the new file was not copied to the destination cluster. This issue is resolved as part of this fix.

Cloudera Bug: OPSAPS-34066

Proper fix to handle parcel activation falsely succeeding when host health is failing

Fixed an issue related to the first run failing when some of the hosts are in bad health. This fix involves adding an extra stop to wait for hosts to report correct the parcel version before proceeding.

Cloudera Bug: OPSAPS-33766

Deferred cgroup removal might not be waiting at all

This fix provides some wait time before the SCM agent attempts to remove control groups (cgroups) again after initial failure.

Cloudera Bug: OPSAPS-32331

SMON should only fall back to Oozie's instrumentation endpoint if the metrics endpoint gives a 503

SMON blindly fell back to Oozie's instrumentation endpoint if it had any problems with the metrics endpoint. Now, SMON falls back only if the metrics endpoint returns a 503 error code, which is the only case it could possibly have success with instrumentation endpoint.

Cloudera Bug: OPSAPS-31988

Support bundle missing client configuration deployment logs

Client configuration deployment logs are now collected as a part of diagnostics support bundle, which were unintentionally removed in Cloudera Manager 5.3.

Cloudera Bug: OPSAPS-31377

Use concurrency option when uploading the sharelib to be faster

You can increase the number of threads used to upload the Oozie sharelib to HDFS. This significantly reduces the time it takes for this operation. This feature is available starting in CDH 5.9.0.

Cloudera Bug: OPSAPS-31078

Host inspector incorrectly warns about kernel version "2.6.32-504.16.2"

Before this fix, the host inspector incorrectly warned about kernel version "2.6.32-504.16.2" as "non-recommended"

Cloudera Bug: OPSAPS-35170

Agent orphan cleanup removes process dir from in flight process

With this fix, the preventative steps in TSB-181 are no longer required.

Cloudera Bug: OPSAPS-35294

Rolling upgrade fails if services stopped

Rolling upgrade will no longer fail if services are stopped.

Cloudera Bug: OPSAPS-34322

Redact Content from Flume configuration

Enabled redaction of sensitive information from Flume configuration.

Cloudera Bug: OPSAPS-34506

Redaction of sensitive information from diagnostic bundles at creation time

Cloudera Manager is designed to transmit certain diagnostic data (or *bundles*) to Cloudera. The Cloudera Support team uses diagnostic bundles to reproduce, debug, and address technical issues for customers. Cloudera support discovered that potentially sensitive data might be included in diagnostic bundles and transmitted to Cloudera. Notwithstanding any possible transmission, such sensitive data is not used by Cloudera for any purpose. Cloudera has taken the following actions:

- Modified Cloudera Manager so that known sensitive data is redacted from the bundles before transmission to Cloudera.
- Updated Cloudera CDH components to remove logging and output of known potentially sensitive properties and configurations.

See [Redaction of Sensitive Information from Diagnostic Bundles](#).

Cloudera Manager 5 Release Notes

Cloudera Bug: OPSAPS-34621

Huge memory leak if the HDFS File Browser UI remains open for an extended time

If the HDFS File Browser were kept open for long periods of time, it could cause a memory leak that would cause Cloudera Manager to crash. This fix addresses the issue.

Cloudera Bug: OPSAPS-35168

Evaluate impact of disabling second-level cache on performance at scale

Disabled second level hibernate cache, which was causing the cache to become "stale" under load. There is no significant difference in Cloudera Manager performance after disabling the cache.

Cloudera Bug: OPSAPS-34010

Express wizard failing to install ZooKeeper

Installing CDH using Packages on Debian 8.2 through the Cloudera Manager wizard could install the incorrect CDH binaries, leading to failures in starting the cluster. Certain packages are available in both the default repository and the Cloudera repository, which could result in the wizard using the wrong one. The Cloudera Manager wizard now ensures the correct binaries are selected. Workaround: Use the parcels format or manually install the correct packages.

Cloudera Bug: OPSAPS-36100

Empty archive files should not be created for impala role diagnostics

Support Bundles started collecting Impala minidumps bundles in Cloudera Manager 5.8.0 and higher. It was generating an empty TAR file if no bundles were found. With this fix, no empty TAR files are generated.

Cloudera Bug: OPSAPS-34887

Cloudera Manager no longer shows "this step is expected to fail" when enabling HDFS-HA

Cloudera Manager now shows the correct label when performing HA.

Cloudera Bug: OPSAPS-30777

If total_space_bytes is very large, heartbeats fail

Fixed an issue where the heartbeat fails with a host that has a mount point backed by cloud storage, such as AWS.

Cloudera Bug: OPSAPS-35742

Enabling cgroups requires "impala" group even if there is no Impala service

Fixed the presumption that any default-named group exists in hosts. YARN cgroup containers do not require the Impala user/group to be present in hosts.

Cloudera Bug: OPSAPS-35242

OOMKiller script not works for Impala Catalog

Fixed a bug where OutOfMemory errors in the Catalog Server might lead to killing multiple Java processes, including other roles on the same host.

Cloudera Bug: OPSAPS-33991

Issues Fixed in Cloudera Manager 5.8.5

Privilege Escalation in Cloudera Manager

Fixed an issue where a read-only Cloudera Manager user can discover the usernames of other users and elevate the privileges of another user. A user cannot elevate their own privilege.

Cloudera Bug: OPSAPS-34124

Use whitelist to only check known good parcels before first run

During the execution of First-Run command, the command only waits for the Cloudera Manager Agent to detect CDH parcels and does not wait for any other activated parcels. Previously, it was waiting for agents to detect all the activated parcels.

Cloudera Bug: OPSAPS-38304

HBase configuration is supplied for Hive when HBase service is not selected

Fixed an issue where Cloudera Manager provided configuration for the `hive-site.xml` file even if the HBase Service setting in Hive is not selected, which could cause unnecessary errors when Hive-on-Spark attempts to connect to HBase. Cloudera Manager now correctly emits the HBase-related configuration in the `hive-site.xml` only when Hive is dependent on HBase.

This issue has been fixed in the following releases:

- CM5.8.5 and higher maintenance releases,
- CM5.9.2 and higher maintenance releases,
- CM5.10.2 and higher maintenance releases,
- CM5.11.1 and higher maintenance releases

When you upgrade Cloudera Manager from an older version to one of the fixed versions, many services may have stale configurations. You must restart all of the stale services in order for them to receive the updated configuration in the `hive-site.xml` file.

Cloudera Bug: OPSAPS-39021

getRoles filtering is not working with Cloudera Manager API version 11 and higher

Fixed a bug that broke backwards compatibility with Cloudera Manager API version 11 (introduced with the Cloudera Manager 5.6 release) for the following endpoint:

```
/api//clusters/clusterName/services/serviceName/roles?filter=type==ROLE_TYPE
```

Cloudera Bug: OPSAPS-40184

Placement rules in DRP UI should provide option to not create sub-pools

Nested User Pools (except existingSecondaryGroup) now support `create=true|false` as a checkbox in the UI.

Cloudera Bug: OPSAPS-39625

Low watermark value for Memstore Flush default is incorrect

For CDH versions 5.8 and higher, the **Low Watermark for Memstore Flush** configuration parameter is associated with the HBase parameter `hbase.regionserver.global.memstore.lowerLimit`.

This value represents the fullness threshold of the memstore as a percentage of memstore capacity. The default value for this parameter was incorrectly set too low at `.38`. This can cause severe under utilization of the memstore.

The default has been corrected to be `.95`. When upgrading to a version of Cloudera Manager with this fix, if the value was previously set to the old default of `.38`, it will automatically be increased to the new default, which may cause Cloudera Manager to mark your HBase service as having a stale configuration, requiring a restart.

Additionally, if an existing **Low Watermark for Memstore Flush** configuration parameter has a value $\leq .9$, it will be flagged as a configuration warning.

Fixed in: Cloudera Manager 5.12, 5.11.1, 5.10.2, 5.9.3, 5.8.5

Cloudera Bug: OPSAPS-38468

Issues Fixed in Cloudera Manager 5.8.4

CM server error - JsonMappingException

Fixed an issue where a race condition can happen when getting the log links for unhealthy roles and then causes a null pointer exception for the Cloudera Manager homepage.

Cloudera Bug: OPSAPS-32804

Dollar sign (\$) char in password results in IllegalArgumentException

Using special characters such as the dollar sign (\$) for passwords and other user-supplied values with CSD-based services (those added to Cloudera Manager with custom service descriptors) raised exceptions. This has been resolved.

Cloudera Bug: OPSAPS-37542

Error when distributing parcels : No such torrent

Parcel distribution can fail, returning the error message:

```
Error when distributing to host: No such torrent:parcel_name.torrent
```

Workaround: Remove the file `/opt/cloudera/parcel-cache/parcel_name.torrent` from the host.

Cloudera Bug: OPSAPS-37183

Hive Replication Command should update copy the Serde properties correctly

Hive Replication now replicates the Serde Properties and also copies the corresponding HDFS file.

Cloudera Bug: OPSAPS-34354

Hive Replication fails when Impala is SSL enabled but Hadoop services are not

Fixes a bug where some Hadoop services did not recognize the Impala SSL configuration, which caused Hive replications to fail when SSL is enabled. To enable SSL, see: [Configuring TLS/SSL for HDFS, YARN and MapReduce](#).

This support also ensures that connection to Impala is successful when SSL is enabled even if Kerberos is not enabled.

Fixed in: Cloudera Manager 5.10.1, 5.9.2, 5.8.4

Cloudera Bug: OPSAPS-38700

Excess load on Cloudera Manager database

Fixed an issue where the load on the Cloudera Manager database increases significantly on clusters with 100-200 nodes or more and when there are host-level configuration overrides. This has been resolved by using a smaller default value for the `default_batch_fetch_size` property in the hibernate configuration.

Fixed in: Cloudera Manager 5.10.0, 5.9.2, 5.8.4, 5.7.6

Cloudera Bug: OPSAPS-37985

Ability to configure the length of Impala queries that SMON keeps in memory

Lengthy queries (such as Impala DDLs) sometimes overloaded the SMON (service monitor) heap, resulting in out-of-memory exceptions (OOM). With this release, queries are limited to 10k characters by default. This setting can be adjusted using the safety valve mechanism (Advanced Configuration Snippet) in Cloudera Manager.

Cloudera Bug: OPSAPS-37053

Issues Fixed in Cloudera Manager 5.8.3

YARN historical reports by user shows pool-user entity

When Cloudera Manager manages multiple clusters, there is no per user tracking for historical applications and queries across clusters. Instead, **Historical Applications by User** and **Historical Queries by User** show applications and queries per user and pool. (A pool is associated with a specific cluster.)

Cloudera Bug: OPSAPS-34986.

If `total_space_bytes` is very large, heartbeats fail

Fixes heartbeat failure with a host that has a mount point backed by cloud storage, such as AWS.

Cloudera Bug: OPSAPS-36144.

OPSAPS-29327 Add config for `hive.metastore.server.max.message.size`

You can configure `hive.meetastore.max.message.size` using **Max Message Size for Hive MetaStore**. The default setting is 100MB. This can cause staleness during a Cloudera Manager upgrade.

Cloudera Bug: OPSAPS-34098.

CM blocks from HS2 enabling both LDAP and Kerberos authentication

HS2 now supports LDAP and Kerberos authentication on the same instance for CDH 5.7.0 and higher.

Cloudera Bug: OPSAPS-34310.

Backport OPSAPS-36100 Debian 8.2 packages `cdh` install fix to C5.8

Installing CDH using Packages on Debian 8.2 through the Cloudera Manager wizard could install the incorrect CDH binaries, leading to failures in starting the cluster. Certain packages are available in both the default repository and the Cloudera repository, which caused the installer to get confused and pick the wrong one. Cloudera Manager's wizard now ensures the correct binaries are selected.

Cloudera Bug: OPSAPS-36971.

Impala `load_catalog_in_background` config should set to "false" by default in CM

Changed default value for `load_catalog_in_background` to false. This fixes catalog scalability issues observed in Impala since v2.2.

Cloudera Bug: OPSAPS-35869.

[api] Latest CM Java client does not work on older CM

Fixed an issue in Cloudera Manager client versions 5.8, 5.8.1, and 5.8.2 where `ApiCollectDiagnosticDataArguments` was incompatible with Cloudera Manager versions lower than 5.8.

Cloudera Bug: OPSAPS-36315

Empty archive files should not be created for Impala role diagnostics

Support Bundles started collecting Impala minidump bundles with Cloudera Manager 5.8 and higher. If no bundles were found, Cloudera Manager generated an empty TAR file. With this fix, no empty TAR files are generated.

Cloudera Bug: OPSAPS-34887

OOMKiller script not works for Impala Catalog

Fixed a bug where `OutOfMemory` errors in the Catalog Server could lead to killing multiple java processes including other roles on the same host.

Cloudera Bug: OPSAPS-33991

Cloudera Manager 5 Release Notes

Sentry Upgrade command is not retrieable

If the Sentry Upgrade command fails, you can now retry the command.

Cloudera Bug: OPSAPS-35365

Support non-public schemas

Cloudera Manager used to work with only 'public' schema in PostgreSQL. With this change, Cloudera Manager supports custom schema names. Cloudera Manager relies on the schema search path in PostgreSQL, which can be set for a user, database, and so on. <https://www.postgresql.org/docs/current/static/ddl-schemas.html>

Cloudera Bug: OPSAPS-33752

Issues Fixed in Cloudera Manager 5.8.2

Improve advanced configuration snippet redaction to encompass cloud provider credentials and other access tokens

The redaction of potentially sensitive parameters in advanced configuration snippets is extended to those commonly used by the Azure Data Lake.

Cloudera Bug: OPSAPS-29523

If files excluded by exclusion filters are renamed, they are not replicated

Customers using Incremental HDFS replication had an issue where if an excluded file (through exclusion filters) is renamed to an 'included' file, the new file is still not copied to the destination cluster. This issue is resolved as part of this fix.

Cloudera Bug: OPSAPS-34066

Hive Replication Command should update copy the Serde properties correctly

Hive Replication now replicates the Serde Properties and also copies the corresponding HDFS file.

Cloudera Bug: OPSAPS-34354

Increase default Solrd watchdog Timeout value

Solr server initialization can take up to 60 secs to complete. During this time interval, Solr server does not respond to the solrd watchdog requests. This can result in solrd watchdog terminating the Solr server process. The default timeout duration for watchdog is increased to 70 secs.

Cloudera Bug: OPSAPS-34320

Add service changes YARN settings

In Cloudera Manager 5.7, adding any new service to your cluster can cause the YARN setting for `mapreduce.job.reduces` to change unexpectedly. Adding a service no longer causes this problem.

Cloudera Bug: OPSAPS-34151

OPSAPS-29327 Add config for `hive.metastore.server.max.message.size`

Hive Metastore max message size can be configured now using **Max Message Size for Hive MetaStore**. It defaults to 100MB. It can cause staleness for customers on Cloudera Manager upgrade.

Cloudera Bug: OPSAPS-34098

XSS in Kerberos activation

In lower releases, there was an XSS vulnerability on the Kerberos page. This is now fixed.

Cloudera Bug: OPSAPS-33883

Upload deployment.json fails when it contains replication info

While trying to migrate Cloudera Manager using deployment.json file with existing Hive replication, schedules used to fail. This has been fixed in this release.

Cloudera Bug: OPSAPS-33927

Cloudera Manager set cataloged default jvm memory to 4G may cause oom on upgrade to Cloudera Manager 5.7+

After upgrading to 5.7 or later, customers could see a reduced Java heap maximum on Impala Catalog Server, due to a change in its default value. Upgrading from Cloudera Manager < 5.7 to Cloudera Manager 5.8.2 no longer sees any effective change in the Impala Catalog Server Java Heap size.

Cloudera Bug: OPSAPS-34039

Oozie first run fails with custom principals

In 5.8.1, Oozie first run fails if using kerberos custom principals. This is now fixed.

Cloudera Bug: OPSAPS-35822

Hive Replication shows "Dry Run" incorrectly

The earlier known issue that running Hive Replication shows "Dry Run" in status message is fixed now.

Cloudera Bug: OPSAPS-33206

HDFS Snapshot policy is selecting unhealthy host to run on

Policy for selecting a role to run HDFS Snapshot command: We select a non-decommissioned host that is in Active status. Also, hosts in maintenance mode have a lower priority than hosts in active state.

Cloudera Bug: OPSAPS-33144

Cluster export fails when service configuration is invalid

The export cluster template code path was failing because of stale configuration in the Cloudera Manager database. Having a stale configuration in the database is possible. This could happen when configurations are deprecated on older CDH releases.

Cloudera Bug: OPSAPS-35586

Agent orphan cleanup removes process dir from in flight process

With this fix, the preventative steps in TSB-181 are no longer required.

Cloudera Bug: OPSAPS-35294

Fix CatalogServiceClient to handle TLS connections to catalogd for UDF replication

When Impala uses SSL, Cloudera supports TLS Connection to Catalog Server. Customers are able to enable replication for any Impala UDFs/Metadata (in Hive Replication).

Cloudera Bug: OPSAPS-34801

Redact Content from Flume config

Enabled redacting sensitive information from Flume configuration.

Cloudera Bug: OPSAPS-34506

Host inspector incorrectly warns about kernel version "2.6.32-504.16.2"

Host inspector incorrectly warns about kernel version "2.6.32-504.16.2" as "non-recommended."

Cloudera Bug: OPSAPS-35170

Cloudera Manager 5 Release Notes

[Impala Breakpad script does not convert exponentials into decimals and leads to errors](#)

Impala Breakpad script failure that happens when trying to collect more than 10 MB of dumps from a single role is fixed.

Cloudera Bug: OPSAPS-34971

[Oozie points to older sharelib even after running sharelib install command](#)

After an "Install Oozie Share Lib" action, the Oozie service is informed that there is a new shared lib installed. This eliminates the need for a separate manual restart.

Cloudera Bug: OPSAPS-26825

Issues Fixed in Cloudera Manager 5.8.1

[CDH upgrade from 5.7.x to 5.8.x fails when Sentry gateway role is enabled](#)

If the Sentry Gateway role is configured on any hosts of a CDH 5.7.x cluster, the upgrade process to CDH 5.8.x fails.

Upgrades to CDH 5.8.x now complete successfully.

Cloudera Bug: OPSAPS-35356

Issues Fixed in Cloudera Manager 5.8.0

[Changes to yarn.nodemanager.remote-app-log-dir not picked up from gateway](#)

YARN log aggregation did not work when `yarn.nodemanager.remote-app-log-dir` was configured to a non-default location. Now this value is emitted in YARN client configuration, ensuring clients logs go to the proper location.

This fix causes staleness.

Cloudera Bug: OPSAPS-29422

[MapReduce2 Counter Limits on Cloudera Manager's YARN page](#)

Fixed an issue where the MapReduce2 property `mapreduce.job.counters.max` was not included in the configuration for the JobHistory Server, which could cause jobs to fail if there were too many counters. This might happen despite increasing the limit configured in Cloudera Manager. The property is now included in the JobHistory Server configuration, in addition to the related property `mapreduce.job.counters.groups.max`.

This fix causes staleness.

Cloudera Bug: OPSAPS-26705

Cloudera Bug: OPSAPS-26705

[Kafka MirrorMaker unable to start due to KAFKA_HOME not being set](#)

Kafka MirrorMaker would not start when Kafka is installed using packages. This occurred because `KAFKA_HOME` was not set to the correct default when starting MirrorMaker. This issue affected Cloudera Manager 5.4.0 and higher with Kafka 1.4.0 and higher.

Cloudera Bug: OPSAPS-33293

[Kafka unable to start due to misconfigured security.inter.broker.protocol when Kerberos is enabled](#)

Kafka would not start when Kerberos is enabled and the default value of `security.inter.broker.protocol` was not changed. This occurred because Kafka tried to use the same port for `SASL_PLAINTEXT` and `PLAINTEXT`. By default, Cloudera Manager now infers the protocol based on the security settings.

This issue affected Cloudera Manager 5.5.2 and higher with Kafka 2.0.0 and higher.

Upgrading to Cloudera Manager 5.7.1 or higher upgrades currently configured values to `INFERRED` unless SSL/TLS is enabled and the values are currently either `PLAINTEXT` or `SASL_PLAINTEXT`. This does not cause any change in behavior.

Cloudera Bug: OPSAPS-31744

Child commands for deleting or adding a nameservice show stack trace

In an existing HDFS deployment with high availability, when you try to add or delete a nameservice and attempt to view the progress of the child commands, a stack trace is triggered if some of the child commands have not yet run. This fix eliminates the stack trace and informs you that the child commands have not yet been run.

Cloudera Bug: OPSAPS-33383

Setting owner of a file in Isilon fails

On Isilon systems, the owner that the file is being changed to must be present on the system. In general cases, the user is not present, so this command fails with an error message suggesting that the user is not part of the supergroup. This fix addresses the issue by not failing the command.

Cloudera Bug: OPSAPS-33145

Handle drop-recreate partition efficiently

With this change, properties of entities (database, table, partition, index) are not updated by default. You can choose to update properties by setting `REPLICATE_PARAMETERS=true` in **Hive Replication Environment Advanced Configuration Snippet (Safety Valve)**.

Cloudera Bug: OPSAPS-33954

HiveReplicationCmdArgs.update is not accessible using Cloudera Manager

In Hive replication, you can choose to update one or more of the entities `INDICES`, `PARAMETERS`, `PARTITIONS`, and `PRIVILEGES` adding the following instruction in **Hive Replication Environment Advanced Configuration Snippet (Safety Valve)**.

```
PROPERTIES_TO_UPDATE=INDICES , PARAMETERS , PARTITIONS , PRIVILEGES
```

Cloudera Bug: OPSAPS-33953

Operation log directory should be configurable and monitored in Cloudera Manager

HiveServer has two new properties for configuration of operation logging.

Property	Default
Enable HiveServer2 Operations Logging	true
HiveServer2 Operations Log Directory	<code>/var/logs/hive/operation_logs</code>

Cloudera Bug: OPSAPS-29483

Kerberos should use non-person objects when creating principals in Active Directory

You can now configure Active Directory account properties. You can use custom values for `objectClasses` to configure accounts, including non-person objects.

Cloudera Bug: OPSAPS-30301

Changes to `yarn.nodemanager.remote-app-log-dir` not picked up from gateway

YARN log aggregation did not work when `yarn.nodemanager.remote-app-log-dir` was configured to a non-default location. Now this value is emitted in YARN client configuration, ensuring clients logs go to the proper location.

Cloudera Bug: OPSAPS-29422

Key Trustee KMS should use round-robin configuration when Key Trustee server uses High Availability

If the Key Trustee server is configured with High Availability, the Key Management Service needs to use round-robin DNS.

Cloudera Bug: OPSAPS-29221

AD_ACCOUNT_PREFIX should not be required

Active Directory Account Prefix was improperly implemented as a required configuration for Security/ Kerberos. The use of this configuration is now optional.

Cloudera Bug: OPSAPS-29196

Proper fix to handle parcel activation falsely succeeding when host health is failing

Fixed an issue related to the first run failing when some of the hosts are in bad health. This fix involves adding an extra stop to wait for hosts to report the correct parcel version before proceeding.

Cloudera Bug: OPSAPS-33766

Document Kerberos + Isilon support in Disaster Recovery

BDR is now supported on Isilon (including on clusters secured with Kerberos).

Cloudera Bug: OPSAPS-33758

Separation of authentication and authorization coprocessor configs in HBase

HBase Secure Bulkload is now enabled for all CDH5.5 and higher clusters, regardless of whether Kerberos is enabled. Also fixed related issue where clusters with authentication (kerberos) but not authorization failed in Hbase-related MapReduce jobs.

Cloudera Bug: OPSAPS-33657

Cloudera Manager - BDR UI - Generate replication diagnostics data. Failed due to java.lang.NullPointerException.

Trying to collect diagnostic data for a Hive replication schedule caused a Java stack trace to be shown on the page. This fix shows an error message instead of displaying a Java stack trace.

Cloudera Bug: OPSAPS-33438

Diagnostic collection fails on a failed BDR job

Failed BDR jobs no longer report errors on "Collect diagnostic data" actions.

Cloudera Bug: OPSAPS-33474

Redact Advanced Configuration Snippets in the UI that contain secrets

Cloudera Manager 5.8 supports redaction of Advanced Configuration Snippet parameters in UI configuration pages. Redaction is based on matching keywords defined as *sensitive*, and detected within the contents of the Advanced Configuration Snippet text. Users who can edit the parameter still see the sensitive words, but users without edit privileges see only the redacted contents.

Cloudera Bug: OPSAPS-33401

Cloudera Manager - snapshot - take snapshot - needs reset of illegal character when fixed

Fixed an issue in the Take Snapshot dialog where the validation of the snapshot name could become "stuck" and prevent execution of the operation.

Cloudera Bug: OPSAPS-33464

BDR UI: Add Search back - It is difficult to find the desired replication schedule in the Cloudera Manager 5.5 UI

The Replications page provides enhanced search functionality to filter schedules by any specified text. Searches occur within any of these schedule fields: HDFS paths, database names and table names.

Cloudera Bug: OPSAPS-33171

Spark standalone does not come up if HDFS is not available

The Spark standalone service now works without an HDFS service. This requires Spark services to show up as stale and require a restart after upgrade to Cloudera Manager 5.7.1 and higher.

Cloudera Bug: OPSAPS-33509

Cloudera Manager - add hdfs nameservice - gets stack trace

In an existing HDFS High Availability setup, when the user tries to add or delete a nameservice and, in the process, attempts to get progress on the child commands, a stack trace is triggered if some of the child commands have not yet run. This fix eliminates the stack trace and informs the user that the child commands have not yet run.

Cloudera Bug: OPSAPS-33383

BDR Replication Schedule page is slow to load

Replication schedules page load performance is improved.

Cloudera Bug: OPSAPS-33125

Replication DistCp - setOwner behavior on Isilon causing failures

On Isilon systems, the owner to which the file is being changed must be present on the system. In general cases, the user is not present, so this command tends to fail with an error message suggesting that the user is not part of the supergroup. This fix addresses the issue by not failing the command.

Cloudera Bug: OPSAPS-33145

Unable to start Hue on cluster that's using Kerberos and Isilon

Hue service can now start with Isilon if Kerberos is enabled.

Cloudera Bug: OPSAPS-27441

Enable database notifications from Hive

Hive now has the property **Enable Stored Notifications in Database**. When set, Hive logs DDL notifications in Hive Metastore.

Cloudera Bug: OPSAPS-32629

The *hbase* user should be whitelisted by default in the list of allowed system users to launch YARN applications

YARN Allowed System Users now includes *hbase* by default. This is helpful when running certain tools for HBase that need to execute MapReduce jobs.

Cloudera Bug: OPSAPS-32631

Allow HDFS Balancer to login with keytab

When running an HDFS rebalance command on a kerberized cluster with a large amount of data, it could would take enough time to complete that authentication would expire and cause an error. Leveraging a new capability in HDFS, the rebalance command is now able to use a keytab file to automatically renew authentication before it expires.

Cloudera Bug: OPSAPS-32372

Cloudera Manager 5 Release Notes

Add support for delete tables and databases deleted on the source

Cloudera Manager now supports deletion of entities from a target Hive database when those entities are deleted from source Hive database during Hive incremental replication.

Cloudera Bug: OPSAPS-32625

Default MapReduce option should be YARN

The default MapReduce service for a new replication schedule is YARN.

Cloudera Bug: OPSAPS-32650

BDR Administrator cannot enable snapshots though the role says it can

BDR Administrator authority message is now more accurate: Create replication schedules and snapshot policies.

Cloudera Bug: OPSAPS-32521

Hive incremental replication enhancements

Cloudera Manager 5.8 includes Hive incremental replication support.

Cloudera Bug: OPSAPS-32366

Cloudera Manager API allows users with read-only privileges to list all Cloudera Manager users

A security bug in the Cloudera Manager API allowed users with read-only permissions to view all the existing users in Cloudera Manager. This issue is now fixed.

Cloudera Bug: OPSAPS-32916

Decide the value for TTL for DbNotifications

You can configure `Time-to-live for Database Notifications` in Hive for notifications present in the `NOTIFICATION_LOG`. The default is 2 days.

Cloudera Bug: OPSAPS-32898

Make the cluster menu sticky

Previously, the Clusters Menu expanded the first cluster by default. As the user expanded or collapsed the accordion, it remembered the configuration for the current session. When the user goes to the services, roles, or host of another cluster, it maintained the configuration of the previously expanded cluster (which might or might not match). In Cloudera Manager 5.7.1 and higher, Cloudera Manager records the last relevant cluster when the user visits a cluster, service, role or host page, and expands that cluster in the Clusters menu by default.

Cloudera Bug: OPSAPS-32850

NPE on Impala Admission Control page if Memory Limit is not set

In lower releases, if the configuration Impala Daemon Memory Limit is not set, the Impala Admission Control page throws a `NullPointerException`. This is now fixed.

Cloudera Bug: OPSAPS-33023

Change Impala service configs for admission control

Enable Impala Admission Control and **Enable Dynamic Resource Pools** are now enabled by default. Customized configuration values are preserved during upgrade.

Cloudera Bug: OPSAPS-32737

Disaster Recovery on Isilon breaks with Kerberos

BDR on Isilon storage is now supported with kerberized clusters.

Cloudera Bug: OPSAPS-32825

Breakpad Crash Reporting for Impala

Support bundles now collect "minidumps" from Impala that help to debug Impala crash issues. As part of this functionality, Cloudera Manager exposes two properties for three roles (Impalad, Catalog Server, Statestore Server).

Property	Description
Breakpad Dump Dir	This determines where the "minidumps" are temporarily available.
Max Breakpad Dump Files	This determines the maximum number of files stored in <code>dump dir</code> (this limits the amount of storage allocated to Impala dump files).

Cloudera Bug: OPSAPS-33050

s3 protocol and scheme should not be pruned during Hive metadata export

With this change, the cloud HDFS path remains as it is after replication.

Cloudera Bug: OPSAPS-32022

YARN `mapreduce.shuffle.max.connections` is a NodeManager setting and not a client setting

`mapreduce.shuffle.max.connections` was emitted to files of YARN clients instead of the NodeManager. It is now correctly emitted only for the NodeManager.

Cloudera Bug: OPSAPS-31772

Kafka unable to start due to listener misconfiguration when Kerberos is enabled

Kafka would not start when Kerberos was enabled and the `security.inter.broker.protocol` default was not changed. This occurred because Kafka would try to use the same port for SASL_PLAINTEXT and PLAINTEXT. By default, Kafka now infers the protocol based on the security settings. This issue affected Cloudera Manager 5.5.2 and higher with Kafka 2.0.0 and higher. Upgrading to Cloudera Manager 5.7.1 and higher upgrades currently configured values to INFERRED unless SSL / TLS is enabled and the values are currently either PLAINTEXT or SASL_PLAINTEXT. This does not cause any change in behavior.

Cloudera Bug: OPSAPS-31744

Chart Builder not showing graphs on IE9

In Cloudera Manager 5.7, charts would not render in the Chart Builder page on IE9. This issue is fixed in Cloudera Manager 5.8.

Cloudera Bug: OPSAPS-31561

Deb 8.2 support for Cloudera Manager

Cloudera Manager is now supported on Debian 8.2.

Cloudera Bug: OPSAPS-31296

ResourceManager would not start if NodeManager were down during the start phase of the restart cycle

All YARN roles are stopped and started together when the service stop or start command is issued with CDH 5.2 and higher. If CDH version is lower than 5.2, the previous behavior of stopping ResourceManagers before NodeManagers and starting them after NodeManagers stays the same.

Cloudera Bug: OPSAPS-30981

Default TLS keystore location for HTTPFS is on non-persistent disk

The default location for the HTTPFS TLS / SSL keystore was `/var/run/hadoop-https/.keystore`, which could be deleted upon machine reboot. Newly created clusters now have an empty default. When upgrading to Cloudera

Manager 5.7.1 or higher, the old value is maintained. There should be no disruption on upgrade, but Cloudera Manager presents a warning that the keystore is in a dangerous location. To fix this problem, move the files to a safe path on the new host, then update the configuration in Cloudera Manager to point to the new path.

Cloudera Bug: OPSAPS-27976

Active Directory principals created without AES 128/256 bit cause job failures if cluster is configured for AES

It is now possible to configure encryption types for Active Directory setups in Cloudera Manager, using a new property on the Kerberos configuration page, **Kerberos Encryption Types**. You can configure only the following 5 encryption types:

- rc4-hmac
- aes128-cts
- aes256-cts
- des-cbc-crc
- des-cbc-md5

Using other values for the encryption type causes validation of this field to fail during upgrade to CDH 5.8.

Cloudera Bug: OPSAPS-27020

MR2 counter limits in the Cloudera Manager YARN page should populate all MR2 config files

The MapReduce2 property `mapreduce.job.counters.max` was not included in the configuration for the JobHistory Server, which could cause jobs to fail if there were too many counters. This might happen despite increasing the limit configured in Cloudera Manager. The property is now included in the JobHistory Server configuration, in addition to the related property `mapreduce.job.counters.groups.max`.

Cloudera Bug: OPSAPS-26705

BDR: support use of a custom principal

Cloudera Manager now supports custom Kerberos principals for BDR.

Cloudera Bug: OPSAPS-24658

Add QueryMonitoring chart to Impala service charts

A new chart on the Impala service page shows query duration for completed Impala queries.

Cloudera Bug: OPSAPS-22909

Add symbols to Cloudera Manager generated Active Directory passwords

Cloudera Manager now allows Active Directory password complexity to be configured using a new security configuration on the UI. The following table lists the default values.

Name	Value
length	12
minLowerCaseLetters	2
minUpperCaseLetters	2
minDigits	2
minSpaces	0
minSpecialChars	0
specialChars	? . ! \$ % ^ * () - _ + = ~

You can modify these values and regenerate Active Directory credentials to create new passwords.

Cloudera Bug: OPSAPS-22853

Regenerate principals should delete from Active Directory, too

It is now possible to regenerate principals for Active Directory setups. This involves deletion of existing accounts and regeneration of new principals. Since some customers might not want to do this through Cloudera Manager, starting with Cloudera Manager 5.8 you can enable the new property `Active Directory Delete Accounts on Credential Regeneration`. Regenerating credentials with this setting enabled automatically deletes existing accounts and completes the regeneration. This setting is disabled by default. If disabled, regeneration of principals throws an error message saying that deletion of accounts is required. Cloudera Manager needs the new configuration to be set in order to delete accounts automatically.

Cloudera Bug: OPSAPS-22182

XSS in Kerberos activation

In lower releases, there was an XSS vulnerability on the Kerberos page. This is now fixed.

Cloudera Bug: OPSAPS-33883

XSS in host addition

In lower releases, there was an XSS vulnerability on the Add Hosts page. This is now fixed.

Cloudera Bug: OPSAPS-33905

XSS in Host Templates

In lower releases, there was an XSS vulnerability on the Host Templates page. This is now fixed.

Cloudera Bug: OPSAPS-33882

Make it more obvious when phone home is off

In lower releases, it was not obvious in the Send Diagnostics dialog whether data would be sent back to Cloudera. The dialog is enhanced to make this information more visible.

Cloudera Bug: OPSAPS-28674

Allow ad hoc sub-pools to be created within existing pools

Cloudera Manager supports the `nestedUserQueue` feature in the UI. You no longer have to use the safety valve to specify `nestedUserQueues`. This means you can make all jobs without specified user queues go to `root.<YOUR_POOL>.<username>` or `root.<YOUR_POOL>.<primaryGroup>` from the UI.

Cloudera Bug: OPSAPS-30156

Add an option to duplicate resource pool configs

Selecting **Clone** lets you create a new pool from the settings of an existing pool.

Cloudera Bug: OPSAPS-32471

Impala Admission Control root pool should have configurable ACLs

Impala Admission Control now supports a global way of editing ACLs.

Cloudera Bug: OPSAPS-33963

Do not create host templates when create cluster wizard is run

During cluster creation, host templates are no longer created automatically.

Cloudera Bug: OPSAPS-32659

Cloudera Manager 5 Release Notes

The modal height is not calculated correctly

In lower releases, the modal dialog is sometimes too tall. This is now fixed.

Cloudera Bug: OPSAPS-33457

new required fields Key Management Server Proxy Group created if there is more than 1 KMS instance

When adding two Key Trustee KMS roles during the initial setup wizard, sometimes these roles were assigned to different groups. The required configuration was set for some but not all of these groups, causing errors. This is now fixed.

Cloudera Bug: OPSAPS-33946

Links should not be shown on print out

When you print pages from Cloudera Manager, the printout no longer displays link URLs.

Cloudera Bug: OPSAPS-33542

Allow customization of ldap account properties for AD-Kerberos setup

The Active Directory account properties `objectClasses` and `accountExpires` are now configurable from the Kerberos Configuration UI page.

Cloudera Bug: OPSAPS-26419

Support bundle has no way to enforce that a certain time is guaranteed to be in bundle

Support bundles can now be collected for a certain time range. Users are also able to get an estimate of the bundle size before collecting the diagnostic data. Only role logs collected as a part of the support bundle are supported by this item in Cloudera Manager 5.8. This item is not available for scheduled support bundles.

Cloudera Bug: OPSAPS-23557

Fix Spark CSD to keep client config files in subdir

The Spark CSD was modified to avoid conflicts with other CSDs that depend on it, and causes the Spark service to show up as stale on upgrade to Cloudera Manager 5.7.1 and higher.

Cloudera Bug: OPSAPS-33511

Cloudera Manager Agent clears out JN data directories that leads to HDFS not restarting

On RHEL 7 class systems, certain configuration actions intended to be executed only once during enablement of HDFS HA might be re-executed when you request a system shutdown or reboot. This can result in data loss. This issue is fixed in Cloudera Manager 5.8, but a hard restart is required on RHEL 7-class systems (including Oracle and CentOS variants) for the fix to take effect. This can be performed after upgrade in a scheduled, rolling manner.

Cloudera Bug: OPSAPS-34171

Hue should not use the embedded sqlite db by default when possible

Hue now has built-in support for PostgreSQL by taking advantage of the system library `python-psycopg2`. In addition to this library, Hue also includes the system libraries listed in the following table.

System	Library
CentOS 5	Not supported
CentOS 6 and 7	postgresql-libs
Ubuntu 10.04	libpq5, python-egenix-mxdatetime and python-central
Ubuntu 12.04 and 14.04	libpq5
SLES 11 sp2	libpq5

Cloudera Bug: OPSAPS-12507

Issues Fixed in Cloudera Manager 5.7.6

The following issues are fixed in Cloudera Manager 5.7.6.

Excess load on Cloudera Manager database

Fixed an issue where the load on the Cloudera Manager database increases significantly on clusters with 100-200 nodes or more and when there are host-level configuration overrides. This has been resolved by using a smaller default value for the `default_batch_fetch_size` property in the hibernate configuration.

Fixed in: Cloudera Manager 5.10.0, 5.9.2, 5.8.4, 5.7.6

Cloudera Bug: OPSAPS-37985

Hive Replication Command should update copy the Serde properties correctly

Hive Replication now replicates the Serde Properties and also copies the corresponding HDFS file.

Cloudera Bug: OPSAPS-34354

Dollar sign (\$) char in password results in IllegalArgumentException

Using special characters such as the dollar sign (\$) for passwords and other user-supplied values with CSD-based services (those added to Cloudera Manager with custom service descriptors) raised exceptions. This has been resolved.

Cloudera Bug: OPSAPS-37542

Hive Replication fails when Impala is SSL enabled but Hadoop services are not

Fixes a bug where some Hadoop services did not recognize the Impala SSL configuration, which caused Hive replications to fail when SSL is enabled. To enable SSL, see: [Configuring TLS/SSL for HDFS, YARN and MapReduce](#).

This support also ensures that connection to Impala is successful when SSL is enabled even if Kerberos is not enabled.

Fixed in: Cloudera Manager 5.10.1, 5.9.2, 5.8.4

Cloudera Bug: OPSAPS-38700

CM server error - JsonMappingException

Fixed an issue where a race condition can happen when getting the log links for unhealthy roles and then causes a null pointer exception for the Cloudera Manager homepage.

Cloudera Bug: OPSAPS-32804

Service Monitor cannot retrieve Navigator Metadata Server metrics when SSL is enabled

Fixed an issue where the Service Monitor cannot retrieve Navigator Metadata Server metrics when SSL is enabled for Navigator Metadata Server.

Fixed in 5.11, 5.10.1, 5.9.2, 5.7.6

Cloudera Bug: OPSAPS-38865

Issues Fixed in Cloudera Manager 5.7.5

The following issues are fixed in Cloudera Manager 5.7.5.

OOMKiller script does not work for Impala Catalog

Fixed a bug where OutOfMemory errors in the Catalog Server might lead to killing multiple Java processes, including other roles on the same host.

Cloudera Bug: OPSAPS-33991

Issues Fixed in Cloudera Manager 5.7.4

The following issues are fixed in Cloudera Manager 5.7.4.

Agent orphan cleanup removes process dir from in flight process

With this fix, the preventative steps in TSB-181 are no longer required.

Cloudera Bug: OPSAPS-35294

YARN historical reports by user shows pool-user entity

When Cloudera Manager manages multiple clusters, there is no per user tracking for historical applications and queries across clusters. Instead, **Historical Applications by User** and **Historical Queries by User** show applications and queries per user and pool. (A pool is associated with a specific cluster.)

Cloudera Bug: OPSAPS-34986

Host inspector incorrectly warns about kernel version "2.6.32-504.16.2"

Host inspector no longer warns that kernel version "2.6.32-504.16.2" as "non-recommended."

Cloudera Bug: OPSAPS-35170

Fix CatalogServiceClient to handle TLS connections to catalogd for UDF replication

When Impala uses SSL, Cloudera supports TLS Connection to Catalog Server. You can enable replication for any Impala UDFs/Metadata (in Hive Replication).

Cloudera Bug: OPSAPS-34801

If total_space_bytes is really big, heartbeats fail

Fixes heartbeat failure with a host that has a mount point backed by cloud storage such as AWS.

Cloudera Bug: OPSAPS-36145

Oozie points to older sharelib even after running sharelib install command

After an "Install Oozie Share Lib" action, the Oozie service is informed that there is a new shared lib installed. This eliminates the need for a separate manual restart.

Cloudera Bug: OPSAPS-26825

Cloudera Manager blocks from HS2 enabling both LDAP and Kerberos auth

HS2 supports LDAP and Kerberos authentication on the same instance for CDH 5.7.0 or higher. Previously, this was considered an error.

Cloudera Bug: OPSAPS-34310

OOMKiller script does not work for Impala Catalog

Fixed a bug where OutOfMemory errors in the Catalog Server could lead to killing multiple java processes, including other roles on the same host.

Cloudera Bug: OPSAPS-33991

Issues Fixed in Cloudera Manager 5.7.2

Unable to start Hue on cluster that's using Kerberos and Isilon

Hue service can now start with Isilon if Kerberos is enabled.

Cloudera Bug: OPSAPS-27441

BDR UI: Add Search back - It is difficult to find the desired replication schedule in the Cloudera Manager 5.5 UI

The Replications page provides enhanced search functionality to filter schedules by any specified text. Searches occur within any of these schedule fields: HDFS paths, database names and table names.

Cloudera Bug: OPSAPS-33171

Handle drop-recreate partition efficiently

With this change, properties of entities (database, table, partition, index) are not updated by default. You can choose to update properties by setting `REPLICATE_PARAMETERS=true` in **Hive Replication Environment Advanced Configuration Snippet (Safety Valve)**.

Cloudera Bug: OPSAPS-33954

Hue `static` directory is browsable

The `static` web directory in Hue is no longer indexed and browsable when the Hue Load Balancer is used.

XSS in Kerberos activation

In lower releases, there was an XSS vulnerability on the Kerberos page. This is now fixed.

Cloudera Bug: OPSAPS-33883

XSS in host addition

In lower releases, there was an XSS vulnerability on the Add Hosts page. This is now fixed.

Cloudera Bug: OPSAPS-33905

XSS in Host Templates

In lower releases, there was an XSS vulnerability on the Host Templates page. This is now fixed.

Cloudera Bug: OPSAPS-33882

Cloudera Manager Agent clears out JN data directories that leads to HDFS not restarting

On RHEL 7 class systems, certain configuration actions intended to be executed only once during enablement of HDFS HA might be re-executed when you request a system shutdown or reboot. This can result in data loss. This issue is fixed in Cloudera Manager 5.7.2, but a hard restart is required on RHEL 7-class systems (including Oracle and CentOS variants) for the fix to take effect. This can be performed after upgrade in a scheduled, rolling manner.

Cloudera Bug: OPSAPS-34171

Separation of authentication and authorization coprocessor configs in HBase

HBase Secure Bulkload is now enabled for all CDH5.5 and higher clusters, regardless of whether Kerberos is enabled. Also fixed related issue where clusters with authentication (kerberos) but not authorization failed in Hbase-related MapReduce jobs.

Cloudera Bug: OPSAPS-33657

Adding a new service changes `mapreduce.job.reduces` setting

Adding a new service in Cloudera Manager no longer changes the YARN setting `mapreduce.job.reduces`.

Cloudera Bug: OPSAPS-34151

BDR: Hive replication status always shows "(Dry Run)"

Running a Hive replication schedule in BDR no longer displays "(Dry Run)" in the status.

BDR Replication Schedule page is slow to load

Replication schedules page load performance is improved.

Cloudera Bug: OPSAPS-33125

HDFS Snapshot running on unavailable nodes

When selecting a host on which to run the **HDFS Snapshot** command, Cloudera Manager now excludes unavailable hosts, such as hosts in maintenance mode or decommissioned hosts.

Cloudera Bug: OPSAPS-33144

Migrating Cloudera Manager using `deployment.json` fails if replication schedules are configured

Migrating Cloudera Manager using `deployment.json` no longer fails if replication schedules are configured.

Cloudera Bug: OPSAPS-33927

Files excluded from replication are not replicated if they are renamed

If a file excluded from replication by an exclusion filter is renamed, it is now replicated properly.

Cloudera Bug: OPSAPS-34066

Issues Fixed in Cloudera Manager 5.7.1

Spark standalone does not come up if HDFS is not available

The Spark standalone service now works without an HDFS service. This requires Spark services to show up as stale and require a restart after upgrade to Cloudera Manager 5.7.1 and higher.

Cloudera Bug: OPSAPS-33509

Kafka unable to start due to listener misconfiguration when Kerberos is enabled

Kafka would not start when Kerberos was enabled and the `security.inter.broker.protocol` default was not changed. This occurred because Kafka would try to use the same port for SASL_PLAINTEXT and PLAINTEXT. By default, Kafka now infers the protocol based on the security settings. This issue affected Cloudera Manager 5.5.2 and higher with Kafka 2.0.0 and higher. Upgrading to Cloudera Manager 5.7.1 and higher upgrades currently configured values to INFERRED unless SSL / TLS is enabled and the values are currently either PLAINTEXT or SASL_PLAINTEXT. This does not cause any change in behavior.

Cloudera Bug: OPSAPS-31744

Default TLS keystore location for HTTPFS is on non-persistent disk

The default location for the HTTPFS TLS / SSL keystore was `/var/run/hadoop-httpfs/.keystore`, which could be deleted upon machine reboot. Newly created clusters now have an empty default. When upgrading to Cloudera Manager 5.7.1 or higher, the old value is maintained. There should be no disruption on upgrade, but Cloudera Manager presents a warning that the keystore is in a dangerous location. To fix this problem, move the files to a safe path on the new host, then update the configuration in Cloudera Manager to point to the new path.

Cloudera Bug: OPSAPS-27976

Fix Spark CSD to keep client config files in subdir

The Spark CSD was modified to avoid conflicts with other CSDs that depend on it, and causes the Spark service to show up as stale on upgrade to Cloudera Manager 5.7.1 and higher.

Cloudera Bug: OPSAPS-33511

Cloudera Manager HDFS usage reports do not include Inode references

Lower versions of Cloudera Manager HDFS usage reports do not include Inode references. As a result, usage reports underreported HDFS directory sizes and data used by users and groups in certain circumstances where HDFS snapshots were used.

Cloudera Bug: OPSAPS-33094

Kafka MirrorMaker unable to start due to KAFKA_HOME not being set

Kafka MirrorMaker would not start when Kafka is installed using packages. This occurred because `KAFKA_HOME` was not set to the correct default when starting MirrorMaker. This issue affected Cloudera Manager 5.4.0 and higher with Kafka 1.4.0 and higher.

Cloudera Bug: OPSAPS-33293

Authentication errors occur due to missing SAML metadata

- If you use SAML for external authentication and were on Cloudera Manager 5.5.0 and higher prior to this upgrade and if you used to notice an error screen while logging out, then this upgrade will fix that issue.
- If you use SAML for external authentication and were on Cloudera Manager 5.5.0 and higher prior to this upgrade and if you did not notice any error screen while logging out, then you will most likely see an error screen while logging out after this upgrade. In order to fix that, you can follow either of these steps:
 1. Update the metadata file in your IdP with the new file from `<Cloudera Manager Server>/saml/metadata`
 2. Change SAML Entity Alias under **Administration** > **Setting "clouderaManager"** to "" and restart Cloudera Manager.

Cloudera Bug: OPSAPS-33088

Child commands for deleting or adding a nameservice show stack trace

In an existing HDFS deployment with high availability, when you try to add or delete a nameservice and attempt to view the progress of the child commands, a stack trace is triggered if some of the child commands have not yet run. This fix eliminates the stack trace and informs you that the child commands have not yet been run.

Cloudera Bug: OPSAPS-33383

HiveServer2 Web UI did not use SSL when Kerberos was enabled

SSL configuration for the HiveServer2 Web UI is now used regardless of whether Kerberos is in use.

Cloudera Bug: OPSAPS-33255

Clusters menu expands to last cluster viewed

Previously, the **Clusters** menu expanded the first cluster by default, and as you expand or collapse the menu, Cloudera Manager remembers that cluster for the session. However, when you go to services, roles, or hosts of another cluster, Cloudera Manager does not remember the other cluster and shows the previously expanded cluster instead.

In release 5.7.1 and higher, Cloudera Manager remembers the last cluster viewed by a user and expands that cluster in the **Clusters** menu by default.

Cloudera Bug: OPSAPS-32850

Impala does not throw null pointer exception when memory limit is not set

If the configuration property **Impala Daemon Memory Limit** was not set, the **Impala Admission Control** page threw a null pointer exception.

Cloudera Bug: OPSAPS-33023

HDFS rolling restart fails after CDH upgrade

In previous Cloudera Manager releases, when one of a pair of highly available NameNodes was down, it was possible for rolling restart or rolling upgrade commands to fail with an error message incorrectly describing the state of the NameNode as "Busy." The error message now correctly identifies the state of the NameNode (typically "Stopped" in this situation).

Cloudera Bug: OPSAPS-33035

The Expand Range option did not work for some charts

The **Expand range to fill all values** option in **Chart Builder** now works for all charts.

Cloudera Bug: OPSAPS-33277

[Kafka unable to start due to misconfigured security.inter.broker.protocol when Kerberos is enabled](#)

Kafka would not start when Kerberos is enabled and the default value of `security.inter.broker.protocol` was not changed. This occurred because Kafka tried to use the same port for `SASL_PLAINTEXT` and `PLAINTEXT`. By default, Cloudera Manager now infers the protocol based on the security settings.

This issue affected Cloudera Manager 5.5.2 and higher with Kafka 2.0.0 and higher.

Upgrading to Cloudera Manager 5.7.1 or higher upgrades currently configured values to `INFERRED` unless SSL/TLS is enabled and the values are currently either `PLAINTEXT` or `SASL_PLAINTEXT`. This does not cause any change in behavior.

Cloudera Bug: OPSAPS-31744

[Oozie JVM heap metrics not reported in Chart Builder for some services](#)

Oozie JVM metrics are now available and display on the role page. They can also be accessed through **Chart Builder** using the `oozie_memory_heap_used` and `oozie_memory_total_max` metrics.

Cloudera Bug: OPSAPS-33006

[Spark CSD modified](#)

The Spark CSD was modified to avoid conflicts with other CSDs that depend on it. This change causes the Spark service to require a restart when upgrading to Cloudera Manager 5.7.1.

Cloudera Bug: OPSAPS-33511

[Poorly formed Advanced Configuration Snippets cause null pointer exception with diagnostic bundles](#)

Certain poorly formed Advanced Configuration Snippets could cause a Null Pointer Exception when uploading diagnostic bundles and setting up a Cloudera Manager peer.

Cloudera Bug: OPSAPS-33378

[Setting owner of a file in Isilon fails](#)

On Isilon systems, the owner that the file is being changed to must be present on the system. In general cases, the user is not present, so this command fails with an error message suggesting that the user is not part of the supergroup. This fix addresses the issue by not failing the command.

Cloudera Bug: OPSAPS-33145

[The Install Oozie ShareLib command is now visible to users with the Configurator role.](#)

Cloudera Bug: OPSAPS-33157

[Default location for TLS Keystore for HTTPFS is nonpersistent](#)

The default location for the HTTPFS TLS / SSL keystore was `/var/run/hadoop-https/.keystore`, which could be deleted when the host reboots. Newly created clusters now have an empty default instead of one that could be deleted. When upgrading to Cloudera Manager 5.7.1 or higher, the old value is maintained, so there is no disruption on upgrade. However, Cloudera Manager warns that the path is in a dangerous location. To fix this problem, move the files to a safe path on that host, and then update the configuration in Cloudera Manager to point to the new path.

Cloudera Bug: OPSAPS-27976

Collecting diagnostic bundle displayed Java stack trace

Collecting a diagnostic bundle for a Hive replication schedule caused a Java stack trace to be shown on the page. This fix shows an error message instead of throwing a Java stack trace.

Cloudera Bug: OPSAPS-33438

Unable to stop Cloudera Manager Agent on SLES 11

Fixes TSB-144.

Running the restart or stop service commands failed to stop the Agent.

Cloudera Bug: OPSAPS-33225

Error creating bean

Occasionally, some users encountered the message `Error creating bean with name 'newServiceHandlerRegistry'` in the Cloudera Manager Admin Console. This issue has been resolved.

Cloudera Bug: OPSAPS-33324

Impala JVM heap size is configurable

The JVM heap size of the Impala catalog server can be configured now using the **Java Heap Size of Catalog Server in Bytes** property. The property defaults to 4 GB, and like all memory parameters may require tuning.

Cloudera Bug: OPSAPS-26483

Issues Fixed in Cloudera Manager 5.7.0

Plain-text passwords sent to users

An authenticated user could request and receive documents that included passwords they were permitted to modify. This information was sent as plain text. Passwords are no longer included in request responses.

Cloudera Bug: OPSAPS-30757

Cloudera Manager forces Solr shutdown before operations complete

When stopping the Solr service, Cloudera Manager would forcibly stop the service after a period of time. This could result in Solr cores not coming up cleanly. Cloudera Manager now waits for all operations to complete on the Solr server before exiting.

Cloudera Bug: OPSAPS-30903

Re-enabling Kerberos fails due to duplicate roles

As part of enabling Kerberos, a role is created. Attempts to re-enable Kerberos failed because the role already existed from when Kerberos was enabled before. When Kerberos is enabled, Cloudera Manager reuses any required roles if they already exist.

Cloudera Bug: OPSAPS-30703

Processes used incorrect HOME variable

CDH service processes and third-party (CSD) processes used an incorrect HOME variable. These processes now run with the HOME environment variable set to the correct home directory based on the process user.

Cloudera Bug: OPSAPS-30610

Stale Kerberos configuration reported after deploying Kerberos client configuration

After making Kerberos configuration changes through **Administration > Settings > Kerberos**, and **Manage krb5.conf** is enabled, the configuration issue 'Cluster has stale Kerberos configuration' might have displayed and might not have disappeared after running **Cluster > cluster_name > Deploy Kerberos Client Configuration**. The deploying Kerberos

client configuration action now waits for pending staleness checks before identifying stale hosts on which to apply Kerberos configuration.

Workaround: Make a different, non-Kerberos edit to a configuration, and save that change. Revert that change immediately afterward.

Cloudera Bug: OPSAPS-30371

Time range settings on report pages are incorrect

Time range settings were incorrect. They are now set correctly on report pages.

Cloudera Bug: OPSAPS-30016

Add Service Wizard fails to set Hive on Spark performance tuning parameters

Automatic configuration for Hive on Spark performance tuning parameters did not run when adding a Hive service to an existing cluster. In the past, these tuning parameters ran when adding a cluster that contained Hive, YARN, and Spark on YARN. The rules now run as long as Hive depends on YARN in both the add service and add cluster wizards.

Cloudera Bug: OPSAPS-25460

Setting empty values for LDAP Base DN and LDAP Domain produces errors

Setting empty values for the LDAP Base DN and LDAP Domain for Impala resulted in errors. These settings are now handled without errors.

Cloudera Bug: OPSAPS-29546

Restart of deleted roles does not wait until deleted roles are identified

Restart and rolling restart commands did not always wait until all deleted roles were identified. As a result, services that had not yet been identified as requiring a restart were not restarted. Restart and rolling restart commands now wait for staleness checks to complete, if the user requests that only stale services be restarted. This avoids a race between staleness computation checking and how services are determined to be stale and thus restarted.

Cloudera Bug: OPSAPS-29190

Miscellaneous problems with the Replication user interface

This release includes several fixes to the replication user interface including the following:

- The **Last Run** column incorrectly sorted dates. Dates are now correctly sorted.
- Collecting diagnostic data for failed Hive Replication commands failed with an error. This data collection now succeeds.
- Finished schedules were shown as running. Schedules now accurately reflect their state.
- Scheduled time was incorrectly translated between browser time and server time. Scheduled time is now correctly translated.
- The **Actions** menu would be disabled while a replication schedule was running, which blocked changing future runs configurations. The **Actions** menu is now enabled for replication schedules with commands running.

Cloudera Bug: OPSAPS-29759

Misleading error occurs while deploying client configuration

The error "There is already a pending command on this entity" was shown if a command was in process for a service and an attempt is made to run the "Deploy client configuration" command. This error is no longer shown.

Cloudera Bug: OPSAPS-31461

Gathering task progress produced null pointer exceptions

Gathering progress information on a variety of tasks could result in null pointer exceptions. For example, this could occur when decommissioning a host. Progress information is now gathered as expected.

Cloudera Bug: OPSAPS-32347

User interface elements are hidden when windows are resized

Some windows could be resized so buttons were not visible. For example, this could happen with the Create Replication dialog box. The user interface now handles resizing as expected.

Cloudera Bug: OPSAPS-32060

HDFS data transfers uses 3DES despite configuration

The configuration information for using the AES/CTR/NoPadding cipher suite for HDFS data transfers was incomplete. As a result, traffic was encrypted with the much slower 3DES algorithm. The correct configuration is now included with the HDFS client.

Cloudera Bug: OPSAPS-31607

Solr keystore passwords are no longer presented in clear text

Because of Tomcat restrictions, Solr keystore passwords were sent as clear text on the machines running the service. They are now redacted.

Cloudera Bug: CDH-33626.

Removing a host from a cluster removes all Kerberos client configuration

Removing a host from a cluster automatically removed all Kerberos client configuration from any other hosts still in the cluster. Now, when one host is removed from a cluster, any Kerberos client configuration on other hosts is unaffected.

Cloudera Bug: OPSAPS-29796

Hive replication import fails to include some information

During the Hive replication import phase, schema and location information was not consistently populated. This information is now populated as expected.

Cloudera Bug: OPSAPS-31332

Restarts attempt to deploy client configuration when action is not supported

After upgrading a parcel, restart or rolling restart steps automatically attempted to deploy the client configuration, even if that was not supported by the cluster configuration. Now, client configurations are deployed only if the cluster configuration supports this action.

Cloudera Bug: OPSAPS-31483

Some pages load very slowly

Some pages, such as the **All Recent Commands** page, may take over 30 seconds to load. This process has been optimized, reducing load times.

Cloudera Bug: OPSAPS-26772

kt_renewer not automatically created when Hue is added to a Kerberos-enabled cluster

When the Hue service was added to a Kerberos-enabled cluster, a single corresponding kt_renewer was not created. The process of adding the Hue service now includes checking if a keytab renewer is required, and creating one if it is.

Cloudera Bug: OPSAPS-31795

Some jobs never run due to memory configuration

It was possible to configure the YARN NodeManager with an amount of available memory less than the amount of memory available to the YARN container. In such a case, a job might never find a NodeManager that meets the memory requirements. The system now ensures that at least one YARN container is configured with an equal or greater amount of memory than the YARN NodeManager value.

Cloudera Bug: OPSAPS-22584

Replication schedule API not compatible with older versions

Clients using the version 10 replication schedule API did not work as expected with instances of Cloudera Manager using version 11 of the API. This meant that clients from Cloudera Manager 5.4.0 and lower did not work as expected with servers running Cloudera Manager 5.5.0 and higher. Clients and servers using these different API versions now function as expected.

Cloudera Bug: OPSAPS-32504

Issues Fixed in Cloudera Manager 5.6.2

Issues Fixed in Cloudera Manager 5.6.1

Scheme and location not filled in consistently during Hive replication import

In previous releases, Hive replication import phase did not consistently fill in scheme and location information. This information is now filled in as expected.

Cloudera Bug: OPSAPS-31522

Cloudera Manager HDFS usage reports do not include Inode references

Lower versions of Cloudera Manager HDFS usage reports do not include Inode references. As a result, usage reports underreported HDFS directory sizes and data used by users and groups in certain circumstances where HDFS snapshots were used.

Cloudera Bug: OPSAPS-33094

Kafka MirrorMaker unable to start due to KAFKA_HOME not being set

Kafka MirrorMaker would not start when Kafka is installed using packages. This occurred because `KAFKA_HOME` was not set to the correct default when starting MirrorMaker. This issue affected Cloudera Manager 5.4.0 and higher with Kafka 1.4.0 and higher.

Cloudera Bug: OPSAPS-33293

kt_renewer not automatically created when Hue is added to a Kerberos-enabled cluster

When the Hue service was added to a Kerberos-enabled cluster, a single corresponding `kt_renewer` was not created. The process of adding the Hue service now includes checking if a keytab renewer is required, and creating one if it is.

Cloudera Bug: OPSAPS-31795

Authentication errors occur due to missing SAML metadata

- If you use SAML for external authentication and were on Cloudera Manager 5.5.0 and higher prior to this upgrade and if you used to notice an error screen while logging out, then this upgrade will fix that issue.
- If you use SAML for external authentication and were on Cloudera Manager 5.5.0 and higher prior to this upgrade and if you did not notice any error screen while logging out, then you will most likely see an error screen while logging out after this upgrade. In order to fix that, you can follow either of these steps:
 1. Update the metadata file in your IdP with the new file from `<Cloudera Manager Server>/saml/metadata`
 2. Change SAML Entity Alias under **Administration** > **Setting "clouderaManager"** to " and restart Cloudera Manager.

Cloudera Bug: OPSAPS-33088

Oozie JVM heap metrics not reported in Chart Builder for some services

Oozie JVM metrics are now available and display on the role page. They can also be accessed through **Chart Builder** using the `oozie_memory_heap_used` and `oozie_memory_total_max` metrics.

Cloudera Bug: OPSAPS-33006

Issues Fixed in Cloudera Manager 5.6.0

CDH upgrade fails if the GPL Extras parcel is in use

CDH and GPL Extras parcel versions must match exactly. When upgrading CDH, by default Cloudera Manager validates the dependency between the new CDH parcel and existing GPL Extras parcel. Since the dependency is not satisfied, the check returns an error and the upgrade fails.

Cloudera Bug: OPSAPS-26436

Workaround:

1. Deactivate parcel dependency checking:
 - a. Select **Administration > Settings**.
 - b. Search for **Validate Parcel Relations**.
 - c. Deselect the checkbox.
 - d. Click **Save Changes** to commit the changes.
2. Deactivate the GPL Extras parcel.
3. Download, distribute, and activate the GPL Extras parcel that matches the CDH upgrade version.
4. Upgrade CDH.
5. Reactivate parcel dependency checking.

Issues Fixed in Cloudera Manager 5.5.6

Separation of authentication and authorization coprocessor configs in HBase

HBase Secure Bulkload is now enabled for all CDH5.5 and higher clusters, regardless of whether Kerberos is enabled. Also fixed a related issue where clusters with authentication (Kerberos) but not authorization failed in HBase-related MapReduce jobs.

Cloudera Bug: OPSAPS-33657

Replications page: sort by schedule ID as default

In the Replications page, the table is now sorted by ID as a default.

Cloudera Bug: OPSAPS-33212

Hive Replication shows "Dry Run" incorrectly

Fixes an issue where running Hive Replication shows "Dry Run" in status message.

Cloudera Bug: OPSAPS-33206

HDFS Snapshot policy is selecting unhealthy host to run on

When selecting a role to run HDFS Snapshot command, Cloudera Manager now selects a non-decommissioned host in Active status. Also, hosts in maintenance mode have a lower priority than the ones in active state.

Cloudera Bug: OPSAPS-33144

Cloudera Manager no longer shows "this step is expected to fail" when enabling HDFS-HA

When enabling HDFS-HA, if the NameNode directories are empty, Cloudera Manager reports this as an error. In lower versions, it returned a message saying the failure is expected. Cloudera Manager now shows the correct message when performing HA.

Cloudera Bug: OPSAPS-30777

Cloudera Manager 5 Release Notes

Configuration for all services is marked stale during upgrade

When you upgrade to version 5.5.6 of Cloudera Manager, the client configuration for all services is marked stale. From the Cluster menu, select Deploy Client Configuration to redeploy the client configuration.

Cloudera Bug: OPSAPS-36234

Issues Fixed in Cloudera Manager 5.5.4

Cluster provisioning fails

In some cases, provisioning of a cluster may fail at the start of the process. This does not happen in all cases and is mainly noticed on RHEL 6 and especially when some hosts are reporting bad health.

Releases affected: 5.5.0-5.5.3, 5.6.0-5.6.1, 5.7.0

Releases containing the fix: 5.5.4, 5.7.1

For releases containing the fix, parcel activation and first run command now completes as expected, even when some hosts report bad health.

This issue is fixed in Cloudera Manager 5.5.4 and 5.7.1 and higher.

Cloudera bug: OPSAPS-33564

Scheme and location not filled in consistently during Hive replication import

In previous releases, Hive replication import phase did not consistently fill in scheme and location information. This information is now filled in as expected.

Cloudera Bug: OPSAPS-31522

Kafka MirrorMaker unable to start due to KAFKA_HOME not being set

Kafka MirrorMaker would not start when Kafka is installed using packages. This occurred because `KAFKA_HOME` was not set to the correct default when starting MirrorMaker. This issue affected Cloudera Manager 5.4.0 and higher with Kafka 1.4.0 and higher.

Cloudera Bug: OPSAPS-33293

Cloudera Manager HDFS usage reports do not include Inode references

Lower versions of Cloudera Manager HDFS usage reports do not include Inode references. As a result, usage reports underreported HDFS directory sizes and data used by users and groups in certain circumstances where HDFS snapshots were used.

Cloudera Bug: OPSAPS-33094

New Hive tables fail to replicate when Sentry Sync is enabled

When Sentry Sync is enabled, new Hive tables failed to replicate. Replication now occurs as expected.

Cloudera Bug: OPSAPS-33065

kt_renewer not automatically created when Hue is added to a Kerberos-enabled cluster

When the Hue service was added to a Kerberos-enabled cluster, a single corresponding `kt_renewer` was not created. The process of adding the Hue service now includes checking if a keytab renewer is required, and creating one if it is.

Cloudera Bug: OPSAPS-31795

Authentication errors occur due to missing SAML metadata

- If you use SAML for external authentication and were on Cloudera Manager 5.5.0 and higher prior to this upgrade and if you used to notice an error screen while logging out, then this upgrade will fix that issue.

- If you use SAML for external authentication and were on Cloudera Manager 5.5.0 and higher prior to this upgrade and if you did not notice any error screen while logging out, then you will most likely see an error screen while logging out after this upgrade. In order to fix that, you can follow either of these steps:
 1. Update the metadata file in your IdP with the new file from <Cloudera Manager Server>/saml/metadata
 2. Change SAML Entity Alias under **Administration** > **Setting "clouderaManager"** to "" and restart Cloudera Manager.

Cloudera Bug: OPSAPS-33088

HDFS rolling restart fails after CDH upgrade

In previous Cloudera Manager releases, when one of a pair of highly available NameNodes was down, it was possible for rolling restart or rolling upgrade commands to fail with an error message incorrectly describing the state of the NameNode as "Busy." The error message now correctly identifies the state of the NameNode (typically "Stopped" in this situation).

Cloudera Bug: OPSAPS-33035

Oozie JVM heap metrics not reported in Chart Builder for some services

Oozie JVM metrics are now available and display on the role page. They can also be accessed through **Chart Builder** using the `oozie_memory_heap_used` and `oozie_memory_total_max` metrics.

Cloudera Bug: OPSAPS-33006

Issues Fixed in Cloudera Manager 5.5.3

[Users using external LDAP authentication with no local Cloudera Manager user role explicitly set may default to the read-only role when upgrading to Cloudera Manager 5.5.2](#)

When upgrading to Cloudera Manager 5.5.2, customers who have non-read-only roles configured through LDAP, and have not explicitly set Cloudera Manager local roles, may lose their Cloudera Manager privileges set by LDAP.

Releases affected: Cloudera Manager 5.5.2

Users affected: Customers who use LDAP for Cloudera Manager user authorization and have upgraded Cloudera Manager from a version lower than 5.5.0 to Cloudera Manager 5.5.2. For example:

- **May be affected:** Install Cloudera Manager 5.3 -> Upgrade to Cloudera Manager 5.5.1 -> Upgrade to Cloudera Manager 5.5.2
- **Unaffected:** Install Cloudera Manager 5.5.1 -> Upgrade to Cloudera Manager 5.5.2

Users not affected:

- Customers who installed Cloudera Manager 5.5.0 and higher and upgraded.
- Customers who use Cloudera Manager local role authorization, regardless of upgrade path and version.

Severity: High

Action required: If you have upgraded to Cloudera Manager 5.5.2 and cannot log in with proper permissions, do the following:

1. Resolve any conflicting user authorization permissions between LDAP and Cloudera Manager local permissions.
2. Contact Cloudera Support for further instructions if you cannot resolve conflicting LDAP and Cloudera Manager user permissions.

If you have not yet upgraded to Cloudera Manager 5.5.2, and are using LDAP user authorization:

1. Before upgrading, resolve any conflicting user authorization permissions between LDAP and Cloudera Manager local permissions.
2. Upgrade to Cloudera Manager 5.5.3 or higher.

Issues Fixed in Cloudera Manager 5.5.2



Note: If you are upgrading from a previous version of Cloudera Manager, Cloudera recommends that you upgrade to version 5.5.3 or higher.

Oozie and HttpFS keystore passwords are no longer presented in clear text

Because of Tomcat restrictions, the Oozie and HttpFS keystore passwords were sent as clear text on the machines running the services. They are now hidden.

Cloudera Bug: OPSAPS-25522

Cross-site scripting vulnerability using malformed strings in the Parcel Remote URLs list

An attacker could set a malformed string in the Parcel Remote URLs list in the database and trigger the attack when a user accesses the Administration Settings page. This attack is now prevented.

Cloudera Bug: OPSAPS-30590

Starting/stopping roles for Flume instance succeeds but displays nothing in popup

In Cloudera Manager 5.5.0, running a Flume start/stop service command would succeed, but display an empty popup. This is now fixed.

Cloudera Bug: OPSAPS-31075

Role process commands missing stderr and stdout in command details

In Cloudera Manager 5.5, certain commands did not show links to `stderr` or `stdout` in the Cloudera Manager UI even if they were executed recently. These could still be found in `/var/run/cloudera-scm-agent/process/` on that host. In Cloudera Manager 5.5.2 `stderr` and `stdout` should appear as they did before.

Note that links to `stderr` and `stdout` for commands may disappear if another command is run on that role. This is expected. The logs can still be found in `/var/run/cloudera-scm-agent/process/` on that host.

Cloudera Bug: OPSAPS-30671

Updating the Hive NameNode location multiple times could lead to data corruption

Multiple updates to the Hive NameNode location could cause Hive Metastore database corruption. Issuing the same command multiple times no longer produces problems.

Cloudera Bug: OPSAPS-30155

Cloudera Manager skips NameNode logs in the diagnostic bundle

Scheduled diagnostic bundles did not include recent role logs. Diagnostic bundles collected manually (not scheduled) worked as expected. Now, scheduled diagnostic bundles include the latest NameNode logs.

Cloudera Bug: OPSAPS-30787

Kafka 2.0 fails to deploy on large clusters as `reserved.broker.max.id` defaults to 1000

Large Kafka clusters would not start when Cloudera Manager-generated broker IDs exceeded the value set by `reserved.broker.max.id`. The default value of `broker.id.generation.enable` has now been set to `false` to disable the `reserved.broker.max.id` configuration property and avoid collisions.

Cloudera Bug: OPSAPS-31128

Cloudera Manager fails to propagate HBase coprocessors to the gateway nodes

Cloudera Manager does not propagate HBase coprocessors to the gateway nodes. As a result, tools that depend on the HBase security subsystem, such as the `loadIncrementalHFiles` tool, do not use security features, even in secure environments.

Cloudera Bug: OPSAPS-28196

Workaround: Add the following properties to the **HBase Client Advanced Configuration Snippet (Safety Valve) for hbase-site.xml** and restart all HBase clients:

```
<property>
  <name>hbase.coprocessor.region.classes</name>
  <value>org.apache.hadoop.hbase.security.token.TokenProvider,org.apache.hadoop.hbase.security.access.SecureBulkLoadEndpoint</value>
</property>
```

HDFS nameservices API call returns incorrect HA role status

The HDFS nameservices API returned incorrect active/standby status information for HA roles. API calls made about roles that were active might return standby status and roles that were in standby status might return active status. Information about host statuses is now accurate.

Cloudera Bug: OPSAPS-29337

Hive requires the `hive` principal for HiveServer2 host as well as load balancer

An issue with HiveServer2 missing its principal and keytab when Hive is load-balanced has been fixed.

Cloudera Bug: OPSAPS-29881

More descriptive error message for service trying to start on a decommissioned host

Cloudera Manager now displays a more descriptive error message when it skips the "Start" command because all roles are started or decommissioned or on a decommissioned host.

Cloudera Bug: OPSAPS-25549

UI shows repeated errors when loading replications page

Issue fixed with loading the replication page when a previous replication command fails without launching a MapReduce job.

Cloudera Bug: OPSAPS-28789

Allow option for external users to be assigned roles in the local database

In Cloudera Manager 5.5, role assignments for external users was disabled, which caused upgrade issues. The fix rolls back the change but instead of using a union approach, implements the following precedence rules:

- If a user is assigned a role in Cloudera Manager, this local role is used.
- Otherwise, a user's LDAP group association determines the user role.

Cloudera Bug: OPSAPS-31181

Clean up usercache directories on migration from unsecure to secure mode

Fixes an issue that led to YARN jobs failing after migration from unsecure to secure mode.

Cloudera Bug: OPSAPS-28378

Spark REST API does not work when parcels are used

The REST API for retrieving data from a live Spark UI or from the Spark History Server has been fixed.

Cloudera Bug: OPSAPS-30919

In secure clusters, DataNode fails to start when `dfs.data.transfer.protection` is set and DataNode ports are changed to unprivileged ports

Before this fix, the only way to run the DataNode on unprivileged ports (port number > 1024) in a Kerberized cluster with DataNode Data Transfer Protection enabled, was to use single-user mode. Now this configuration works for both regular and single-user mode installs.

Both Hadoop SSL and DataNode Data Transfer Protection are still required for unprivileged DataNode ports to work in a Kerberized cluster. This configuration is supported only in CDH 5.2 and higher.

Cloudera Bug: OPSAPS-30819

New validation warning for non-recommended secure DataNode configurations in CDH 5.2 and higher

On a Kerberos-enabled cluster running CDH 5.2 and higher, there are two recommended DataNode configurations. Use SASL/TLS with DataNode Data Transfer Protection enabled to encrypt the connection, or use only privileged ports to communicate. The supported combinations of HDFS configuration properties follow:

- **Security through SASL/TLS (preferred):**
 - DataNode Data Transfer Protection - Enabled
 - Hadoop TLS/SSL - Enabled
 - DataNode Transceiver Port - Non-privileged (that is, port number \geq 1024)
 - Secure DataNode Web UI Port (TLS/SSL) - Non-privileged (that is, port number \geq 1024)
- **Security through privileged ports:**
 - DataNode Data Transfer Protection - Disabled
 - Hadoop TLS/SSL - Disabled
 - DataNode Transceiver Port - Privileged (that is, port number < 1024)
 - DataNode HTTP Web UI Port - Privileged (that is, port number < 1024)

Any configuration other than these results in a validation warning or error from Cloudera Manager. In particular, the following configuration, which is allowed by HDFS but is not recommended, results in a (dismissible) validation warning:

- **(Not Recommended) Security without enabling DataNode Data Transfer Protection:**
 - DataNode Data Transfer Protection - Disabled
 - Hadoop TLS/SSL - Enabled
 - DataNode Transceiver Port - Privileged (that is, port number < 1024)
 - Secure DataNode Web UI Port (TLS/SSL) - Non-privileged (that is, port number \geq 1024)

All configurations other than the three listed result in Cloudera Manager displaying a validation error.

Cloudera Bug: OPSAPS-31266

Sensitive environment parameters not redacted for CSDs

Passwords in environment variables for CSDs are now redacted.

Cloudera Bug: OPSAPS-29915

Update Apache Commons Collections library in Cloudera Manager due to major security vulnerability

The Apache Commons Collections library has been upgraded to 3.2.2 to fix a critical security vulnerability.

Cloudera Bug: OPSAPS-30428

Remove plaintext keystore password from `/api/v6/cm/config`

With the addition of the JVM parameter, `\-Dcom.cloudera.api.redaction=true`, sensitive configuration values are redacted from the API.

Cloudera Bug: OPSAPS-26677

JVM parameter to redact passwords now redacts the password salt and hash

When API redaction is turned on using the JVM argument `-Dcom.cloudera.api.redaction=true`, it also redacts the user's `pwHash` and `pwSalt` values. Passwords for Cloudera Manager Peers are also redacted.

Cloudera Bug: OPSAPS-31194

Oozie keystore and truststore passwords now redacted

Oozie's Java keystore and truststore passwords are no longer sent in clear text on the command line.

Cloudera Bug: OPSAPS-25523

Several Replication UI fixes

- The Replication UI **Last Run** column now sorts correctly based on dates.
- Collecting diagnostic data for failed Hive Replication commands no longer fails.
- Finished schedules are no longer shown as running when the user is watching the page.
- Scheduled time now accurately translates between browser time and server time.
- Actions menu is now enabled for replication schedules with commands running. Previously, it would be disabled while a replication schedule was running, which blocked changing the configuration for future runs.

Cloudera Bug: OPSAPS-29750

Fix Java detection

Java detection in the "Components" view for a host was fixed to account for Java versions installed using symlinks in `/usr/java` (such as `/usr/java/default`). A similar fix was made to the host inspector's Java detection logic.

Cloudera Bug: OPSAPS-31105

When Host Monitor is stopped, cluster/services status in Cloudera Manager API returns Good instead of N/A or Unknown

If the Host Monitor is down, the service details page will still be able to present non-host related health status.

Cloudera Bug: OPSAPS-30337

Issues Fixed in Cloudera Manager 5.5.1

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera ("Cloudera Products"), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled "Arbitrary remote code execution with InvokerTransformer" and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Issues Fixed in Cloudera Manager 5.5.0

[Setting the HBase WAL provider to the default in Cloudera Manager causes the RegionServer process to fail to start](#)

Setting the HBase WAL provider to the default using Cloudera manager erroneously sets the value of `hbase.wal.provider` to `default`, when it should be set to `defaultProvider`. This causes the RegionServer process to fail to start.

Cloudera Bug: OPSAPS-29598

Workaround: Do not use Cloudera Manager to set the WAL provider to the default. Instead, add the following properties to the **RegionServer Advanced Configuration Snippet (Safety Valve) for hbase-site.xml** and restart all HBase clients.

```
<property>
  <name>hbase.wal.provider</name>
  <value>defaultProvider</value>
</property>
```

[Incorrect path format causes access permission failure](#)

The file path format used in sequence file `ing was incorrect. In replications involving a large number of files and when the **Delete Policy** for the replication is set to **Delete to trash** or **Delete Permanently**, distcp uses the local file system to save the intermediate result of sequence file sorting. An incorrect path format causes access permission failure.

Cloudera Bug: OPSAPS-28859

[Out-of-memory exception for Hive replication](#)

Hive replication throws an out-of-memory exception when exporting a table with a large number of partitions.

Cloudera Bug: OPSAPS-29598

[Incorrect value emitted into hbase-site.xml](#)

When configuring an HBase WAL provider with the "HBase Default" option, Cloudera Manager emits an incorrect value into `hbase-site.xml` and HBase reports an error.

Cloudera Bug: OPSAPS-29598

[Failed replication no longer fail silently](#)

When a replication is attempted between secure and insecure clusters, the replication reports an error and fails silently.

Cloudera Bug: OPSAPS-28380

[The Replication Schedules page displays error after failed replication](#)

A dialog box that shows a Java exception displays in the Replication Schedules page after a failed replication and the page is inaccessible.

Cloudera Bug: OPSAPS-28465

[Cloudera Manager now allows you to use '/' in cluster names](#)

In previous versions, this resulted in problems during replication because '/' was treated as an URL path.

Cloudera Bug: OPSAPS-26138

HDFS replication fails because of improper snapshot directory handling

Replications fail when converting a snapshot path to a regular path when the grandparent of the source directory is snapshottable.

Cloudera Bug: OPSAPS-27232

Renewal time limits and lifetime limits are removed for Kerberos tickets

Snapshots that take longer than 30 minutes failed because the Kerberos tickets for the snapshots expired too soon . The renewal and lifetime limits have been replaced with the system default lifetime and renewal limits, instead of 30 minutes.

Cloudera Bug: OPSAPS-28189

Replication schedules fails to catch configuration errors during creation

The schedule configuration is not validated when it is created, which causes errors during replication.

Cloudera Bug: OPSAPS-11047

Audit events now include the schedule ID

The replication ID has been added to the schedule creation, update, and deletion audit events.

Cloudera Bug: OPSAPS-11534

OutOfMemory error when running HDFS replication

HDFS replication fails with an OutOfMemory exception (in `stderr.log`) when an HDFS replication job replicates a large number of files.

Cloudera Bug: OPSAPS-28858

Hive replication does not preserve user and group names for the database directory

File permissions of the database directory that maps to a Hive database are not preserved.

Cloudera Bug: OPSAPS-24742

Exception in configuration history after changing host configuration

An `IllegalStateException` is reported in the History and Rollback page after changing host configuration.

Cloudera Bug: OPSAPS-28795

Host Inspector no longer expects the impala user to be in the hdfs group

To conform to security best practices, the impala user should not be in the hdfs group.

Cloudera Bug: OPSAPS-28129

The DataNode Refresh Data Directories command fails on secure clusters

DataNode refresh failed on Kerberized clusters.

Cloudera Bug: OPSAPS-27068

Topology should only include hosts with Hadoop daemons

A topology map is created for the services HDFS, YARN, and MapReduce. In releases earlier than 5.5.0, hosts with only gateway roles for these services were also added to the topology map. This might erroneously mark many configurations as stale, requiring restarts or refreshes to resolve the staleness.

Cloudera Bug: OPSAPS-24109

Agent `host_id` property is erroneously set to hostname

Previously, attempting to set the `listening_hostname` property in the Agent `config.ini` file (which is not normally necessary) changed the Agent's host ID to use this hostname, instead of the normal value.

Cloudera Bug: OPSAPS-27991

`JAVA_HOME` override setting does not affect component list

The component list for a host is not affected by a custom `JAVA_HOME` setting.

Cloudera Bug: OPSAPS-19290

Monitoring performance issues on large, busy clusters

A number of fixes have been made to improve Service Monitor and Host Monitor performance, particularly problems manifesting as large, regular garbage-collection pauses, on large, busy clusters.

Cloudera Bug: OPSAPS-26703

`JAVA_HOME` does not get passed to Agents

Client configuration deployment fails to locate Java in the custom `JAVA_HOME` environment variable, if that is specified via host configuration.

Cloudera Bug: OPSAPS-27048

Job History Server Retaining Logs in Secure Clusters

The Job History Server fails to delete old logs from HDFS on secure clusters, with the error "Failed to specify server's Kerberos principal name."

Cloudera Bug: OPSAPS-30157

On Kerberized clusters, incorrect values reported for DataNode process metrics

The DataNode process is incorrectly monitored on Kerberized clusters.

Cloudera Bug: OPSAPS-28676

During `kt_renewer` kinit call authentication may fail

On Kerberized clusters, Cloudera Manager might periodically report monitoring failures due to authentication errors when attempting to communicate with various role web servers to collect metrics. This is due to synchronization issues between the Agent's calls to kinit and attempts to perform Kerberos authentication with the role web servers. This has been fixed by adding logic to retry requests for metrics in the Agent a number of times on authentication failures. These retries will be logged but should not result in health failures and alerts.

Cloudera Bug: OPSAPS-28469

Client configurations are incorrectly marked as stale when a host is rebooted

Rebooting a host can cause the client configurations for gateway roles on that host to be marked as stale even though they are up to date.

Cloudera Bug: OPSAPS-28568

Rolling restart fails when inheriting inappropriate JVM properties

Rolling restart inherits custom HBase RegionServer JVM properties and can fail when those properties are inappropriate for non-daemon JVMs.

Cloudera Bug: OPSAPS-28353

Clusters can fail if Java 1.6 is installed

`JAVA_HOME` is set to Java 1.6 if installed, even if version 1.8 is also installed.

Cloudera Bug: OPSAPS-28888

Start should clearly indicate when it fails due to a missing parcel

Previously a missing parcel was logged only in the Agent log. The Cloudera Manager Admin Console now indicates that a required parcel is missing when starting a role or deploying client configurations.

Cloudera Bug: OPSAPS-25589

Client configuration deployment timeout set too large for large number of hosts

The deploy client configuration timeout value was set to according to the number of hosts. This caused a problem when there are large number of hosts. The tasks to deploy client configurations are run concurrently so there was no need to wait that long. The timeout value was changed to a fixed value.

Cloudera Bug: OPSAPS-29464

Two HBase metrics charts display "No Data"

HBase IPC metrics were not collected on CDH 5.4 and higher HBase due to a metric name change.

Cloudera Bug: OPSAPS-28171

Agent operating system detection logic fixed on updated Oracle Enterprise Linux 6.x

An Oracle Enterprise Linux 6.x update to its `python-libs` unexpectedly changed the output of the Agent's operating system detection logic, which caused problems with parcel distribution and other issues.

Cloudera Bug: OPSAPS-28193

ZooKeeper `jute.maxbuffer` emitted into configuration instead of JVM arguments

The ZooKeeper `jute.maxbuffer` property was emitted into `zoo.cfg` instead of in the JVM arguments. It is now passed to the JVM through the process environment variable `ZOOKEEPER_SERVER_OPTS` and takes effect correctly.

Cloudera Bug: OPSAPS-28244

High `-Xms` set for Hive clients

`-Xms` is set equal to `-Xmx` for all Hive clients which causes the Java runtime to reserve `-Xms` memory even if the client does not need it. This particularly affects machines with low resources. The `-Xms` setting has been removed so that the Java runtime does not assign `-Xmx` memory at the start. Instead, it starts with a much lower Java heap and increases it if needed.

Cloudera Bug: OPSAPS-28972

Service autoconfiguration set Kafka memory to 0

The Kafka service was initially configured with 0 memory because the autoconfiguration rules did not respect specified units. This is now correctly set to a value between 50 and 1024 MiB depending on available memory on the host. This issue affected any third-party CSD-based services using memory parameters with units other than "bytes".

Cloudera Bug: OPSAPS-27869

Topics with a period not reporting metrics

Topics with a period (".") in the name would not show metrics in Cloudera Manager.

Sensitive information in Cloudera Manager diagnostic support bundles

Cloudera Manager is designed to transmit certain diagnostic data (or "bundles") to Cloudera. These diagnostic bundles are used by the Cloudera support team to reproduce, debug, and address technical issues for our customers. Cloudera internally discovered a potential vulnerability in this feature, which could cause any sensitive data stored as "advanced configuration snippets (ACS)" (formerly called "safety valves") to be included in diagnostic bundles and transmitted to Cloudera. Notwithstanding any possible transmission, such sensitive data is not used by Cloudera for any purpose.

Cloudera has taken the following actions: (1) Modified Cloudera Manager so that it no longer transmits advanced configuration snippets containing the sensitive data, and (2) Modified Cloudera Manager TLS/SSL configuration to increase the protection level of the encrypted communication.

Cloudera strives to follow and also help establish best practices for the protection of customer information. In this effort, we continually review and improve our security practices, infrastructure, and data-handling policies.

Users affected:

- Users storing sensitive data in advanced configuration snippets

Impact: Possible transmission of sensitive data

CVE: CVE-2015-6495

Immediate Action Required:

- Upgrade Cloudera Manager to one of the following releases: Cloudera Manager 5.5.0, 5.4.6, 5.3.7, 5.2.7, 5.1.6, 5.0.7, 4.8.6

Cloudera Management Service can fail with a large Flume configuration file

When a Flume configuration file is large, calling its Kerberos credentials with regex can cause the Cloudera Management Service to time out and fail. In addition, the Cloudera Manager Server uses 100% of the CPU and the UI hangs.

Cloudera Bug: OPSAPS-29685

Charts built on simple select statements can return partial results

Charts built on queries in the form `<select metric>` for service or role metrics might not filter entities properly and might report hitting the stream limit.

Cloudera Bug: OPSAPS-28716

Issues Fixed in Cloudera Manager 5.4.10

Scheme and location not filled in consistently during Hive replication import

In previous releases, Hive replication import phase did not consistently fill in scheme and location information. This information is now filled in as expected.

Cloudera Bug: OPSAPS-31522

Having `hive.compute.query.using.stats` enabled by default produced incorrect results for some queries that used stats only

By default, `hive.compute.query.using.stats` was enabled. This produced incorrect results for some queries that used stats only. This setting is now disabled by default.

Cloudera Bug: OPSAPS-32332

YARN jobs fail after enabling Kerberos authentication or selecting Always Use Container Executor

After Kerberos security is enabled on a cluster or Always Use Container Executor is selected, YARN jobs failed. This occurred because the contents of any previously existing YARN User Cache directory could not be overridden after security was enabled. YARN jobs now complete as expected after a change in Kerberos security or usage of Container Executor.

Cloudera Bug: OPSAPS-32050

Scheduled diagnostic bundles do not include recent role logs

Diagnostic bundles collected manually included all expected logs, but bundles collected on a schedule did not include role logs. Scheduled diagnostic bundles now include all expected logs.

Cloudera Bug: OPSAPS-30787

Issues Fixed in Cloudera Manager 5.4.9

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera (“Cloudera Products”), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled “Arbitrary remote code execution with InvokerTransformer” and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Cross-site scripting vulnerability using malformed strings in the parcel remote URL list

An attacker could set a malformed string in parameters that consist of a list of strings and trigger the attack when a user accessed the corresponding configuration page in classic layout mode. This attack is now prevented.

Cloudera Bug: OPSAPS-30590

Cross-site scripting vulnerability using malformed host template name

An attacker could set a malformed host template name in the backend database and trigger the attack when a user applies the host template. This attack is now prevented.

Cloudera Bug: OPSAPS-30443

Snapshot policies with names with special characters not handled as expected

Snapshot policies with names containing special characters such as #, \$, ?, or % were not handled as expected. These snapshot policies were not consistently found because the special characters in their names were parsed incorrectly. Snapshot policies with names containing special characters are now handled as expected.

Cloudera Bug: OPSAPS-30426

Kafka MirrorMaker fails to find messages if ZooKeeper root directory is changed

When the ZooKeeper root directory is changed, the corresponding value in ZK_QUORUM that is passed to Kafka MirrorMaker processes is not updated. In that case, MirrorMaker fails to find messages. Changes to ZooKeeper root are now propagated properly, resulting in MirrorMaker finding messages.

Cloudera Bug: OPSAPS-30197

Updating the Hive NameNode location multiple times could lead to data corruption

Multiple updates to the Hive NameNode location could cause Hive metastore database corruption. Issuing the same command multiple times no longer produces problems.

Cloudera Manager 5 Release Notes

Cloudera Bug: OPSAPS-30155

Cloudera Manager monitors subject to excessive garbage-collection workload

The Cloudera Manager Service Monitor and Cloudera Manager Host Monitor wrote aggregate timeseries data in a way that resulted in significant garbage-collection workloads. Writes are now split based on metric threshold counts, resulting in lower garbage-collection loads.

Cloudera Bug: OPSAPS-30058

Kafka MirrorMaker fails to start because of missing settings

Kafka MirrorMaker provided no defaults for Destination Broker List, Topic Whitelist, and Topic Blacklist, and no way to set these values in the wizard. These values can now be set in the wizard when adding a new instance.

Cloudera Bug: OPSAPS-29415

Cloudera Manager dry-run replication history shows unexpected values

BDR replication in dry-run mode should show the number of files that would be copied and the number of bytes those files would comprise if the same job were executed without the dry-run option. Dry-run mode showed the number of replicable files accessed up to a maximum of 1024 files and showed the total number of bytes those files comprise, up to 512 bytes per file.

The results of dry-runs now show the actual number of source files and their composite bytes that would be covered in the replication schedule. These categories are labeled replicable files and replicable bytes.

Cloudera Bug: OPSAPS-30387

Issues Fixed in Cloudera Manager 5.4.8

Clusters can fail if Java 1.6 is installed

JAVA_HOME is set to Java 1.6 if installed even if 1.8 is also installed.

Cloudera Bug: OPSAPS-28888

OutOfMemory error when running HDFS replication

HDFS replication fails with an OutOfMemory exception (in `stderr.log`) when an HDFS replication job replicates a large number of files.

Cloudera Bug: OPSAPS-28858

Cloudera Management Service can fail with a large Flume configuration file

When a Flume configuration file is large, calling its Kerberos credentials with regex can cause the Cloudera Management Service to timeout and fail. In addition, the Cloudera Manager Server uses too much CPU (100%) and the UI hangs.

Cloudera Bug: OPSAPS-29685

Client configurations are incorrectly marked as stale when a host is rebooted

Rebooting a host can cause the client configurations for gateway roles on that host to be marked as stale even though they are actually up to date.

Cloudera Bug: OPSAPS-28568

Issues Fixed in Cloudera Manager 5.4.7

Operating system detection logic for Oracle Enterprise Linux 6 breaks parcel distribution

Fixes the operating system detection logic on updated Oracle Enterprise Linux 6 systems. An Oracle update to its python-libs logic unexpectedly changed the output of the Agent's operating system detection logic, which caused problems with parcel distribution and other issues.

Cloudera Bug: OPSAPS-28193

ZooKeeper `jute.maxbuffer` property emitted into the wrong file

The ZooKeeper `jute.maxbuffer` property is emitted into `zoo.cfg` instead of in the JVM arguments. It is now passed to the JVM through the process environment variable `ZOOKEEPER_SERVER_OPTS` and takes effect correctly.

Cloudera Bug: OPSAPS-28244

Create user API call is allowed for user with insufficient permissions

Using the "create a user" API call, a user who normally could not create users is able to create a read-only user account. The API call now respects the permissions.

Cloudera Bug: OPSAPS-27539

Spark Authentication property is propagated to the wrong client configuration

Beginning with Cloudera Manager 5.4.6, the **Spark Authentication** configuration property is correctly propagated to client configurations.

Cloudera Bug: OPSAPS-27912

Agent hostname in `config.ini` is changed to wrong value

Attempting to set the `listening_hostname` property in the Agent's `config.ini` file (which is not normally necessary) changes the Agent's host ID to use this hostname, instead of the normal value. The host ID is now left unchanged, as expected.

Cloudera Bug: OPSAPS-27991

Cloudera Manager monitors the incorrect process for DataNode

On Kerberized clusters, Cloudera Manager monitors the wrong process as the DataNode. That has been fixed. For customers using Kerberized HDFS, Cloudera Manager reports incorrect statistics in some areas (memory, file descriptor, CPU usage, I/O, and networking, but not HDFS statistics). There is a small impact to health monitoring because of this issue. For customers using the stacks collection feature on a Kerberized DataNode and where `jstack` collection was enabled, this issue kills the parent `jsvc` process of the DataNode and leaves the DataNode up, but causes Cloudera Manager to report the process as dead.

Cloudera Bug: OPSAPS-28677

Issues Fixed in Cloudera Manager 5.4.6

Sensitive Information in Cloudera Manager Diagnostic Support Bundles

Cloudera Manager is designed to transmit certain diagnostic data (or "bundles") to Cloudera. These diagnostic bundles are used by the Cloudera support team to reproduce, debug, and address technical issues for our customers. Cloudera internally discovered a potential vulnerability in this feature, which could cause any sensitive data stored as "advanced configuration snippets (ACS)" (formerly called "safety valves") to be included in diagnostic bundles and transmitted to Cloudera. Notwithstanding any possible transmission, such sensitive data is not used by Cloudera for any purpose.

Cloudera has taken the following actions: (1) modified Cloudera Manager so that it no longer transmits advanced configuration snippets containing the sensitive data, and (2) modified Cloudera Manager TLS/SSL configuration to increase the protection level of the encrypted communication.

Cloudera strives to follow and also help establish best practices for the protection of customer information. In this effort, we continually review and improve our security practices, infrastructure, and data handling policies.

Users affected:

- Users storing sensitive data in advanced configuration snippets

Impact: Possible transmission of sensitive data

CVE: CVE-2015-6495

Immediate Action Required:

- Upgrade Cloudera Manager to one of the following releases: Cloudera Manager 5.4.6, 5.3.7, 5.2.7, 5.1.6, 5.0.7, 4.8.6

Issues Fixed in Cloudera Manager 5.4.5

Cancel Impala Query attempts to connect via TLS/SSL despite TLS/SSL being disabled

In Impala queries, if you select Cancel for any query, you see a small "internal error" at the top of the query list. This occurs because an attempt to connect via TLS/SSL is performed even though Impala does not have TLS/SSL enabled.

Cloudera Bug: OPSAPS-27146;

Cloudera Manager displays a spurious validation warning about the Cloudera Management Service truststore

Cloudera Manager incorrectly warns that Cloudera Management Service daemons will use HTTPS for communication with either Cloudera Manager or CDH services, even if no Cloudera Management Service truststore is in use.

Cloudera Bug: OPSAPS-27473;

Aggregation of Work attributes

Cloudera Manager now correctly aggregates Work attributes such as YARN applications or Impala query duration.

Cloudera Bug: OPSAPS-26493

Hue Solr Indexer

Cloudera Manager now creates the correct configuration required to create a collection.

Cloudera Bug: OPSAPS-22975

Cloudera Manager incorrectly reports "Not finalized" status for rolling upgrade

When performing a rolling upgrade from a version of CDH lower than 5.4 to CDH 5.4, and the HDFS rolling upgrade is finalized, Cloudera Manager incorrectly reports the status as not finalized. This is an error in reporting only and does not affect HDFS functionality.

Cloudera Bug: OPSAPS-27379

Validation errors not visible from service-level configuration pages

Validation errors and warnings were only visible when accessing the individual instance-level configuration pages. This has been fixed.

Cloudera Bug: OPSAPS-27066

Add Role Instances wizard does not work when initialized using the Cloudera Manager search box

You can now start the Add Role Instances wizard by searching for "<service name> Add Role" in the Cloudera Manager Admin Console search box.

Cloudera Bug: OPSAPS-27604

Cloudera Manager now allows you to use '/' in cluster names

In previous versions, this resulted in problems during replication because '/' was treated as an URL path.

Cloudera Bug: OPSAPS-26138

Expose HBase multi-WAL configuration properties in Cloudera Manager

The properties were added and are now being written to the `hbase-site.xml` file.

Cloudera Bug: OPSAPS-27139

Custom Kerberos principals now handled correctly during Solr startup

The Solr custom Kerberos principal is now initialized properly during Solr server startup.

Cloudera Bug: OPSAPS-27320

Added a check in the Upgrade and Kerberos wizards to make sure Spark-standalone is not enabled

Spark Standalone does not work in clusters with Kerberos authentication. Spark on YARN supports Kerberos and is recommended over Spark Standalone. Either disable Kerberos or remove Spark Standalone before upgrading.

Cloudera Bug: OPSAPS-26151

Fixed link for Reports when YARN high availability is enabled

The Reports link in HA mode would result in a 404 error.

Cloudera Bug: OPSAPS-27089

Removed bogus failure when deploying client configuration

Deploying client configuration would sometimes fail because Cloudera Manager could not locate `JAVA_HOME`. This is not a valid failure because deploying client configuration does not require Java.

Cloudera Bug: OPSAPS-27650

Added `core-site.xml` to Sentry's classpath

Previously, `core-site.xml` was only added to Sentry's configuration folder, but not the classpath.

Cloudera Bug: OPSAPS-26617

Improved memory usage in serializing objects and writing them to support bundles

Performance improvements that require less memory were made for the creation of bundles.

Cloudera Bug: OPSAPS-27405

New property to enable suppressing INFO-level log messages from NameNode

You can now use the **NameNode Block State Change Logging Threshold** property to suppress INFO-level block state change log messages from the NameNode.

Cloudera Bug: OPSAPS-26437

Improved advice for clock offset health test

The way the health of the host's NTP daemon is determined was changed recently, which caused some cases where the related health test (host clock offset) failed without a warning. Information on this change was added to Cloudera Manager.

Cloudera Bug: OPSAPS-27278

Cloudera Manager displays warning about using RHEL 6 with Transparent Huge Pages (THP)

The THP algorithm was broken in certain variants of RHEL 6.2 and above. Cloudera Manager now displays a warning if THP is enabled for all RHEL 6 and above.

Cloudera Bug: OPSAPS-27035

Agent gets no logs if the last log4j event is larger than the max-size specified

If the `byte_limit` (`max-size`) specified by Cloudera Manager during log retrieval was smaller than the last log4j event to be collected, the Agent skipped the complete event and return nothing. This behavior was modified so pick the first N bytes (`N = max-size`) are picked from the log4j event and return a partial log4j event.

Cloudera Bug: OPSAPS-26406

Agent log retrieval does not always honor timeouts

Cloudera Manager Agents no longer enter an infinite loop during log retrieval.

Cloudera Bug: OPSAPS-26404

Cloudera Manager Agent missing log messages

The default timeout for displaying log entries (`./logs/search` and `./logs/context`) has been increased to 60 seconds.

Cloudera Bug: OPSAPS-27358

Fixed cross-site scripting vulnerability

A cross-site scripting vulnerability was discovered and fixed in Cloudera Manager.

Cloudera Bug: OPSAPS-27496

New property added for ResourceManager high availability failover

The ZooKeeper session timeout property `yarn.resourcemanager.zk-timeout-ms` was added, and its default value is 1 minute.

Cloudera Bug: OPSAPS-20852

Set maximum value for YARN `mapreduce.jobhistory.max-age-ms` to 10 years

Cloudera Manager would previously display a validation error when the value was greater than 60 days.

Cloudera Bug: OPSAPS-27182

Added warning in upgrade wizard regarding dropped support for symlinks in CDH 5

This fix added a warning about removing HDFS symlinks when upgrading from CDH 4 to CDH 5.

Cloudera Bug: OPSAPS-26665

Refresh Data Directories command no longer fails on secure clusters

The `DataNodeRefreshCommand` now sets `SCM_KERBEROS_PRINCIPAL` in the environment of the command process, which causes `hdfs.sh` to do a `kinit`. Before this change, a manual `kinit` was required.

Cloudera Bug: OPSAPS-27068

New Sentry Synchronization Path Prefixes added in NameNode configuration are not enforced correctly

Any new path prefixes added in the NameNode configuration are not correctly enforced by Sentry. The ACLs are initially set correctly, however they would be reset to the old default after some time interval.

Cloudera Bug: OPSAPS-26141

Workaround: Set the following property in **Sentry Service Advanced Configuration Snippet (Safety Valve)** and **Hive Metastore Server Advanced Configuration Snippet (Safety Valve) for hive-site.xml**:

```
<property>
<name>sentry.hdfs.integration.path.prefixes</name>
<value>/user/hive/warehouse, ADDITIONAL_DATA_PATHS</value>
</property>
```

where `ADDITIONAL_DATA_PATHS` is a comma-separated list of HDFS paths where Hive data will be stored. The value should be the same value as `sentry.authorization-provider.hdfs-path-prefixes` set in the `hdfs-site.xml` on the NameNode.

Fixed NullPointerException on health tests' Details page

The health tests **Details** page threw a NullPointerException because it was referring to a deprecated metric name.

Cloudera Bug: OPSAPS-27065

Improved Service Monitor Canary check to see if HTable is disabled

Without this check, Service Monitor would fail due to too many ZooKeeper connection messages leaking into the Service Monitor log. This resulted in resource and allocation pressures on the Service Monitor.

Cloudera Bug: OPSAPS-27318

Cloudera Manager no longer retains unnecessary references to HTables

Retaining too many unnecessary references to HTable was using up too much memory, especially when working with a large number of tables.

Cloudera Bug: OPSAPS-27319

Sqoop 2 failure in Kerberized clusters fixed

Cloudera Manager was using the wrong authentication package and picking up the wrong configuration properties for Sqoop 2 authentication with Kerberos.

Cloudera Bug: OPSAPS-27297

Fixed Solr server startup error

The Solr server would not start due to insufficient space for the shared memory file.

Cloudera Bug: OPSAPS-27158

Added HBase Canary security configuration properties

Enabling the HBase canary on a secure cluster would fail. The new properties now let Cloudera Manager specify the canary's Kerberos principal and keytab in the `hbase-site.xml` deployed at the RegionServers.

Cloudera Bug: OPSAPS-26468

Issues Fixed in Cloudera Manager 5.4.3

Improve Impala queries coordinator node metrics handling

For Impala queries that returned very few rows, Cloudera Manager could fail to report information such as HDFS I/O metrics on the **Impala Query Monitoring** and **Query Detail** pages. The discrepancy was typically relatively small because those queries often did very little work.

Cloudera Bug: OPSAPS-26531, OPSAPS-26533

Performance issues when changing configurations on HDFS

Fixed a performance issue where HDFS configuration pages responded slowly.

Cloudera Bug: OPSAPS-27171

Typo in Cloudera Manager metrics reference

The word "Concerning" was misspelled in many metrics reference pages.

Cloudera Bug: OPSAPS-26957

Issues with Navigator field `audit_log_max_file_size`

The log4j appender changed from `RollingFileAppender` to `RollingFileWithoutDeleteAppender`.

Cloudera Bug: OPSAPS-26978

Cloudera Manager 5 Release Notes

The Isilon client configuration `core-site.xml` file does not contain proxy users

The parameters are available in the Cloudera Manager Admin Console, but the configurations are not emitted in the `core-site.xml` file.

Cloudera bug: OPSAPS-26816

Solr gateway role should not have a `log4j.properties` advanced configuration snippet

The Solr gateway role does not have a `log4j.properties` file.

Cloudera bug: OPSAPS-26858

The Cloudera Manager Agent `force_start`'s hard stop commands did not set all invariants

This resulted in NPE being reported in Cloudera Manager logs when accessing active and recent command operations.

Cloudera bug: OPSAPS-26864

Configuration staleness icons appear to be enabled for users in read-only role

When moused over, the icons change to a hand indicating that they are active. However, users in the read-only role cannot act on changed configurations.

Cloudera bug: OPSAPS-26897

Setting `yarn.resourcemanager.am.max-retries` throws error

Observed when setting the Maximum Number of Attempts for MapReduce Jobs and then setting ApplicationMaster Maximum Attempts, which also sets `yarn.resourcemanager.am.max-retries`.

Cloudera bug: OPSAPS-26936

Cloudera Manager reports the wrong value for Impala bytes read from cache

Instead of cached bytes it reported the value of short circuit bytes.

Cloudera bug: OPSAPS-26938

Fixed cross-site scripting vulnerabilities

A variety of possible cross-site scripting vulnerabilities have been fixed.

Cloudera bugs: OPSAPS-26765, OPSAPS-26798, OPSAPS-26835, OPSAPS-26836, OPSAPS-26878, OPSAPS-26880, OPSAPS-25959

Location of Number of rows drop-down changed

On pages where multiple rows display, the drop-down menu where users select the number of rows to display on a page now appears at the bottom of all lists.

Cloudera bug: OPSAPS-26956

Minimum allowed value change for YARN property

The **Max Shuffle Connections** property now allows a value of 0, which indicates no limit on the number of connections.

Cloudera bug: OPSAPS-26953

Upgrade error

A bug was fixed that prevented upgrades from CDH 4.7.1 to CDH 5.4.3.

Cloudera bug: OPSAPS-27119

Change to Parcels page

On the Parcels page, the first cluster in the list is now automatically selected by default.

Cloudera bug: OPSAPS-27045

All Password Input Fields do not allow auto complete

All password input fields in Cloudera Manager do not allow auto complete.

Cloudera bug: OPSAPS-26927

TLS Keystore Configuration Error

It is no longer possible to delete the values of the **Path to TLS Keystore File** and **Keystore Password** properties and save them while the **Use TLS Encryption for Admin Console** property is enabled.

Cloudera bug: OPSAPS-26888

Host configuration properties and Agent restart messages

Some host configuration properties no longer incorrectly state that an Agent restart is required.

Cloudera bug: OPSAPS-26808

More detailed error messages for failed role migration

If there is a failure validating the NameNode or JournalNode data directories while migrating roles, Cloudera Manager now displays detailed error information, including error codes.

Cloudera bug: OPSAPS-26803

New property to configure Oozie shared library upload timeout

To prevent timeouts due to slow disks or networks, a new Oozie property, **Oozie Upload ShareLib Command Timeout**, has been added to set the timeout.

Cloudera bug: OPSAPS-26782

New Cluster-Wide Configuration Pages

The following new Cluster-Wide configuration pages have been added:

- Databases
- Local Data Directories
- Local Data Files
- Navigator Settings
- Service Dependencies

To access these pages in Cloudera Manager, select **Cluster > Cluster Name > Configuration**.

Cloudera bug: OPSAPS-26730

Naming of Health Tests

The names of some Health Tests have changed to use consistent capitalization.

Cloudera bug: OPSAPS-26690

Impala Monitoring Queries for Per-node peak memory

Impala queries that report per-node peak memory were incorrect when the value is zero.

Cloudera bug: OPSAPS-26721

Enable Hive on Spark Property

The description of the **Enable Hive on Spark** property has been updated to remind the user that the **Enable Spark on YARN** property must also be selected.

Cloudera bug: OPSAPS-26560

Cloudera Manager 5 Release Notes

Role Trigger property in Flume

Setting a value for the Flume **Role Triggers** property no longer causes validation warnings.

Cloudera bug: OPSAPS-26743

Restart of Service Monitor leaves files that can fill the disk

Restarts of the Service Monitor no longer leave extraneous copies of files that unnecessarily take up disk space.

Cloudera bug: OPSAPS-26660

HiveServer 2 properties omit Java options

Setting any of the following properties no longer causes Java options to be omitted:

- **Allow URIs in Database Policy File**
- **HiveServer2 TLS/SSL Certificate Trust Store File**
- **HiveServer2 TLS/SSL Certificate Trust Store Password**

Cloudera bug: OPSAPS-26657

CDH Parcel distribution reports HTTP 503 errors

Cloudera Manager no longer displays HTTP 503 errors during distribution of the CDH parcel to a large cluster.

Cloudera bug: OPSAPS-26627

Diagnostic bundle reports incorrect status for SELinux

Diagnostic bundles sometimes reported SELinux as disabled when it was enabled. The bundle now reports the correct status.

Cloudera bug: OPSAPS-26608

Hue configuration warnings do not link to correct page

On the Cloudera Manager page that displays Hue configuration issues, the links now take the user to the correct page where the user can correct the configuration.

Cloudera bug: OPSAPS-26600

Date display in Cloudera Manager log viewer

The month and date have been added before the time value in logs displayed in Cloudera Manager.

Cloudera bug: OPSAPS-26599

Disabling Hive Metastore Canary Test

When you disable the Hive Metastore health test by deselecting the **Hive Metastore Canary Health** property, the Hive Canary is now also disabled.

Cloudera bug: OPSAPS-26588

Agent failure when TLS 1.0 is disabled

If TLS 1.0 is disabled, the Agent now tries to negotiate the connection using TLS 1.1 or TLS 1.2.

Cloudera bug: OPSAPS-26584

Slowness when displaying details of a stale configuration

The details page now displays more quickly when a user clicks on the Stale Configuration icon.

Cloudera bug: OPSAPS-26537

Slowness observed when accessing replication page in Cloudera Manager

When you access the replication page in Cloudera Manager, the page responds slowly due to a large number of replication history records. The number of displayed historical records has been changed from 100 to 20.

Cloudera Bug: OPSAPS-26529;

Log Searches for Cloudera Manager Server

Searching the Cloudera Manager Server logs now works as expected.

Cloudera bug: OPSAPS-26470

Failed TLS Configuration and Cloudera Manager Restart

If the TLS configuration has errors, Cloudera Manager now falls back to non-TLS operation when restarting.

Cloudera bug: OPSAPS-27104

New headers added

New headers have been added to Cloudera Manager HTTP response headers to protect against vulnerabilities.

Cloudera bug: OPSAPS-25847

Hive Logging property restored

The `Enable Explain Logging (hive.log.explain.output)` property was removed in an earlier release and is now included in the configurations.

Cloudera bug: OPSAPS-25857

Hive Metastore Update NameNodes Command

A 150 second timeout was removed from the Update Hive Metastore NameNodes command to prevent timeouts on deployments that use Hive extensively.

Cloudera bug: OPSAPS-26056

Kafka Parcel Installation

Cloudera Manager now correctly detects the Kafka version for parcel installation.

Cloudera bug: OPSAPS-26071

Agent restart failure

In a condition where an Agent restart was required due to a Hive configuration change and a subsequent disk failure, the Agent now restarts as expected.

Cloudera bug: OPSAPS-26172

Error message wording

Some Cloudera Manager error messages referred to Cloudera Manager as "CM". These messages now use the full name "Cloudera Manager".

Cloudera bug: OPSAPS-26191

Oozie metrics failures

Retrieval of Oozie metrics sometimes fails due to timeout issues which are now resolved.

Cloudera bug: OPSAPS-26242

NameNode Role Migration Failures

When a NameNode role migration fails due to the destination role data directories being non-empty or having incorrect permissions, you no longer need to complete the migration manually. An error message displays and you can now correct the problem and re-run the command.

Cloudera bug: OPSAPS-26265

AWS S3 HBase configuration property renamed to Amazon S3

Several configuration properties for HBase have been renamed from AWS S3 to Amazon S3, in order to use the correct product name.

Cloudera bug: OPSAPS-23343

NodeManager Host Resources page display for the NodeManager Recovery Directory

The NodeManager Recovery Directory now displays on the NodeManager host resources page.

Cloudera bug: OPSAPS-23916

Host Inspector page now includes link to Show Inspector Results

The Host Inspector page now displays a link to a page that displays detailed results.

Cloudera bug: OPSAPS-24642

Initialization Script Improvements

The Cloudera Manager Agent initialization script now checks correctly for running processes.

Cloudera bug: OPSAPS-24941

Default Value for Hue parameter changed

The default value for the Hue `cherry_server_threads` property has been changed from 10 to 50.

Cloudera bug: OPSAPS-25030

Express Installation Wizard Package Installation Page CDH Version

The Express Installation Wizard Package installation page no longer allows the user to proceed without selecting a CDH version.

Cloudera bug: OPSAPS-25722

Host Component page display

The Host Component page now displays the package version for the KMS Trustee Key Provider.

Cloudera bug: OPSAPS-25692

Installation Wizard hangs during package installation

The Installation Wizard hangs during a CDH package installation and the status displays as "Acquiring Installation Lock". A bug was fixed where the Agent incorrectly failed to release a lock until the Agent is restarted.

Cloudera bug: OPSAPS-22894

Minimum allocation violation not caught by Cloudera Manager

NodeManager did not start because Cloudera Manager did not correctly validate memory and CPU settings against their minimum values.

Cloudera bug: OPSAPS-17157

Impala core dump directories are now configurable

Three new properties that specify the location of core dump directories have been added to the Impala configurations:

- Catalog Server Core Dump Directory
- Impala Daemon Core Dump Directory
- StateStore Core Dump Directory

Cloudera bug: OPSAPS-21743

Typo in Sqoop DB path suffix (SqoopParams.DERBY_SUFFIX)

Sqoop 2 appears to lose data when upgrading to CDH 5.4. This is due to Cloudera Manager erroneously configuring the Derby path with "reposito" instead of "repository". The correct path name is now used.

Cloudera Bug: OPSAPS-26649.

Agent fails when retrieving log files with very long messages

When searching or retrieving large log files using the Agent, the Agent no longer consumes near 100% CPU until it is restarted. This can also happen then the collect host statistics command is issued.

Cloudera bug: OPSAPS-25854

Automated Solr TLS/SSL configuration may fail silently

Cloudera Manager 5.4.1 offers simplified TLS/SSL configuration for Solr. This process uses a `solrctl` command to configure the `urlSchemeSolr` cluster property. The `solrctl` command produces the same results as the Solr REST API call `/solr/admin/collections?action=CLUSTERPROP&name=urlScheme&val=https`. For example, the call might appear as:

```
https://example.com:8983/solr/admin/collections?action=CLUSTERPROP&name=urlScheme&val=https
```

Cloudera Manager automatically executes this command during Solr service startup. If this command fails, the Solr service startup now reports an error.

Cloudera Bug: OPSAPS-26563

Removing the default value of a property fails

For example, when you access the **Automatically Downloaded Parcels** property on the following page: **Home > Administration > Settings** and remove the default `CDH` value, the following error message displays: "Could not find config to delete with template name: `parcel_autodownload_products`". This error has been fixed.

Cloudera bug: OPSAPS-26591

Issues Fixed in Cloudera Manager 5.4.1

distcp default configuration memory settings overwrite MapReduce settings

Replication used for backup and disaster recovery does not correctly set the MapReduce Java options, and you cannot configure them. In release 5.4.1, Cloudera Manager uses the MapReduce gateway configuration to determine the Java options for replication jobs. Replication job settings cannot be configured independently of MapReduce gateway configuration. See [Backup and disaster recovery replication does not set MapReduce Java options](#) on page 2541.

Oozie high availability plug-in is now configured by Cloudera Manager

In CDH 5.4.0, Oozie added a new HA plugin that allows all of the Oozie servers to synchronize their Job ID assignments and prevent collisions. Cloudera Manager 5.4.0 did not configure this new plugin; Cloudera Manager 5.4.1 now does so.

Cloudera Bug: OPSAPS-25778

HDFS read throughput Impala query monitoring property is misleading

The `hbase_bytes_read_per_second` and `hdfs_bytes_read_per_second` Impala query properties have been renamed to `hbase_scanner_average_bytes_read_per_second` and `hdfs_scanner_average_bytes_read_per_second` to more accurately reflect that these properties return the average throughput of the query's HBase and HDFS scanner threads, respectively. The previous names and descriptions indicated that these properties were the query's total HBase and HDFS throughput, which was not accurate.

Cloudera Bug: OPSAPS-26140

Enabling wildcarding in a secure environment causes NameNode to fail to start

In a secure cluster, if you use a wildcard for the NameNode's RPC or HTTP bind address, the NameNode fails to start. For example, `dfs.namenode.http-address` must be a real, routable address and port, not `0.0.0.0.port`. In Cloudera Manager, the "Bind NameNode to Wildcard Address" property must not be enabled. This should affect you only if you are running a secure cluster and your NameNode needs to bind to multiple local addresses.

Cloudera Bug: CDH-9991, OPSAPS-14577

Bug: HDFS-4448

Workaround: Disable the "Bind NameNode to Wildcard Address" property found on the Configuration tab for the NameNode role group.

Support for adding Hue with high availability

The Express and Add Service wizards now allow users to define multiple Hue service roles. If Kerberos is enabled, a co-located KT Renewer role is automatically added for each Hue server row.

Cloudera Bug: OPSAPS-21639

Parameter validation fails with more than one Hue role

When you add a second Hue role to a cluster, the error message "Failed parameter validation" displays.

Cloudera Bug: OPSAPS-25336

Cross-site scripting vulnerabilities

Various cross-site scripting vulnerabilities were fixed.

Cloudera Bug: OPSAPS-25809

Clicking the "Revert to default" icon stores the default value as a user-defined value in the new configuration pages

Cloudera Manager 5.4.1 fixes an issue in which saving an empty configuration value causes the value to be replaced by the default value. The empty value is now saved instead of the default value.

Cloudera Bug: OPSAPS-26340

Spurious validation warning and missing validations when multiple Hue Server roles are present

When multiple Hue Server roles are created for a single Hue Service, Cloudera Manager displays a spurious validation warning for Hue with the label "Failed parameter validation." The Cloudera Manager Server log may also contain exception messages of the form:

```
2015-03-30 17:15:45,077 WARN
ActionablesProvider-0:com.cloudera.cmf.service.ServiceModelValidatorImpl:
Parameter validation failed java.lang.IllegalArgumentException: There is more than one
role with roletype: HUE_SERVER [...] {
```

These messages do not correspond to actual validation warnings and can be ignored. However, some validations normally performed are skipped when this spurious warning is generated, and should be done manually. Specifically, if Hue's authentication mechanism is set to LDAP, the following configuration should be validated:

1. The Hue **LDAP URL** property must be set.

2. For CDH 4.4 and lower, set one (but not both) of the following two Hue properties: **NT Domain** or **LDAP Username Pattern**.
3. For CDH 4.5 and higher, if the Hue property **Use Search Bind Authentication** is selected, exactly one of the two Hue properties **NT Domain** and **LDAP Username Pattern** must be set, as described in step 2 above.

Cloudera Bug: OPSAPS-25336

Logging of command unavailable message improved

When a command is unavailable, the error messages are now more descriptive.

Cloudera Bug: OPSAPS-26384

Client configuration logs no longer deleted by the Agent

If the Agent fails to deploy a new client configuration, the client log file is no longer deleted by the agent. The Agent saves the log file and appends new log entries to the saved log file.

Cloudera Bug: OPSAPS-26148

HDFS role migration requires certain HDFS roles to be running

Before using the Migrate Roles wizard to migrate HDFS roles, you must ensure that the following HDFS roles are running as described:

- A majority of the JournalNodes in the JournalNode quorum must be running. With a quorum size of three JournalNodes, for example, at least two JournalNodes must be running. The JournalNode on the source host need not be running, as long as a majority of all JournalNodes are running.
- When migrating a NameNode and co-located Failover Controller, the other Failover Controller (that is, the one that is not on the source host) must be running. This is true whether or not a co-located JournalNode is being migrated as well, in addition to the NameNode and Failover Controller.
- When migrating a JournalNode by itself, at least one NameNode / Failover Controller co-located pair must be running.

Cloudera Bug: OPSAPS-25870, OPSAPS-25878

HDFS role migration requires automatic failover to be enabled

Migration of HDFS NameNode, JournalNode, and Failover Controller roles through the Migrate Roles wizard is only supported when HDFS automatic failover is enabled. Otherwise, it causes a state in which both NameNodes are in standby mode.

Cloudera Bug: OPSAPS-25806, OPSAPS-25822

HDFS/Hive replication fails when replicating to target cluster that runs CDH 4 and has Kerberos enabled

Cloudera Bug: OPSAPS-25945

Workaround: None.

Issues Fixed in Cloudera Manager 5.4.0

Proxy Configuration in Single User Mode is Fixed

In single user mode, all services are using the same user to proxy other users in an unsecure cluster, which is the user that is running all the CDH processes on the cluster. To restrict that user so that it can proxy other users from only certain hosts and only certain groups, configure the **YARN Proxy User Hosts** and **YARN Proxy User Groups** properties in the HDFS service. The setting here supersedes all other proxy user configurations in single user mode.

Cloudera Bug: OPSAPS-25518

The Parcels page allows access to the patch release notes

Clicking the icon with an "i" in a blue circle next to a parcel shows the release notes.

Cloudera Bug: OPSAPS-25481

Monitoring Fails on Impala Catalog Server with TLS/SSL Enabled

When enabling TLS/SSL for Impala web servers (`webserver_certificate_file`), Cloudera Manager does not emit `use_ssl` in the `cloudera-monitor.properties` file for the Catalog Server. Other services (Impala Daemon and StateStore) are configured correctly. This causes monitoring to fail for the Catalog Server even though it is working as expected.

Cloudera Bug: OPSAPS-25469

Default Value Changed for `hive.exec.reducers.bytes.per.reducer`

To improve performance, the default value for the `hive.exec.reducers.bytes.per.reducer` property has been changed from 1 GB to 64 MB. If this value has been customized, the customized value is retained during an upgrade. If the old default of 1 GB was not changed, the value will be updated to 64MB during an upgrade.

Cloudera Bug: OPSAPS-24883

Issues Fixed in Cloudera Manager 5.3.10

Scheme and location not filled in consistently during Hive replication import

In previous releases, Hive replication import phase did not consistently fill in scheme and location information. This information is now filled in as expected.

Cloudera Bug: OPSAPS-31522

YARN jobs fail after enabling Kerberos authentication or selecting Always Use Container Executor

After Kerberos security is enabled on a cluster or Always Use Container Executor is selected, YARN jobs failed. This occurred because the contents of any previously existing YARN User Cache directory could not be overridden after security was enabled. YARN jobs now complete as expected after a change in Kerberos security or usage of Container Executor.

Cloudera Bug: OPSAPS-32050

Issues Fixed in Cloudera Manager 5.3.9

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera (“Cloudera Products”), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled “Arbitrary remote code execution with InvokerTransformer” and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Cloudera Manager monitors subject to excessive garbage-collection workload

The Cloudera Manager Service Monitor and Cloudera Manager Host Monitor wrote aggregate timeseries data in a way that resulted in significant garbage-collection workloads. Writes are now split based on metric threshold counts, resulting in lower garbage-collection loads.

Cloudera Bug: OPSAPS-30058

Cloudera Manager dry-run replication history shows unexpected values

BDR replication in dry-run mode should show the number of files that would be copied and the number of bytes those files would comprise if the same job were executed without the dry-run option. Dry-run mode showed the number of replicable files accessed up to a maximum of 1024 files and showed the total number of bytes those files comprise, up to 512 bytes per file.

The results of dry-runs now show the actual number of source files and their composite bytes that would be covered in the replication schedule. These categories are labeled replicable files and replicable bytes.

Cloudera Bug: OPSAPS-30387

Updating the Hive NameNode location multiple times could lead to data corruption

Multiple updates to the Hive NameNode location could cause Hive metastore database corruption. Issuing the same command multiple times no longer produces problems.

Cloudera Bug: OPSAPS-30155

Issues Fixed in Cloudera Manager 5.3.8

Rolling Restart fails when inheriting inappropriate JVM properties

Rolling Restart inherits custom HBase RegionServer JVM properties and can fail when those properties are inappropriate for non-daemon JVMs.

Cloudera Bug: OPSAPS-28353

Exception thrown when viewing host configuration change

When ConfigContext is initialized to both cluster and host, it defaults to cluster. If the context is a host, this can cause the ConfigTableUtil class to throw an IllegalStateException.

Cloudera Bug: OPSAPS-28795

Issues Fixed in Cloudera Manager 5.3.7

Sensitive Information in Cloudera Manager Diagnostic Support Bundles

Cloudera Manager is designed to transmit certain diagnostic data (or “bundles”) to Cloudera. These diagnostic bundles are used by the Cloudera support team to reproduce, debug, and address technical issues for our customers. Cloudera internally discovered a potential vulnerability in this feature, which could cause any sensitive data stored as “advanced configuration snippets (ACS)” (formerly called “safety valves”) to be included in diagnostic bundles and transmitted to Cloudera. Notwithstanding any possible transmission, such sensitive data is not used by Cloudera for any purpose.

Cloudera has taken the following actions: (1) modified Cloudera Manager so that it no longer transmits advanced configuration snippets containing the sensitive data, and (2) modified Cloudera Manager TLS/SSL configuration to increase the protection level of the encrypted communication.

Cloudera strives to follow and also help establish best practices for the protection of customer information. In this effort, we continually review and improve our security practices, infrastructure, and data handling policies.

Users affected:

- Users storing sensitive data in advanced configuration snippets

Severity: High

Impact: Possible transmission of sensitive data

CVE: CVE-2015-6495

Immediate Action Required:

- Upgrade Cloudera Manager to one of the following releases: Cloudera Manager 5.4.6, 5.3.7, 5.2.7, 5.1.6, 5.0.7, 4.8.6

Issues Fixed in Cloudera Manager 5.3.6

[Cloudera Manager reports the wrong value for Impala bytes read from cache](#)

Instead of cached bytes, it reported the value of short-circuit bytes.

Cloudera Bug: OPSAPS-26938

[Cancel Impala Query attempts to connect via TLS/SSL despite TLS/SSL being disabled](#)

In Impala queries, if you select Cancel for any query, you get a small "internal error" at the top of the query list. This is because an attempt to connect via TLS/SSL is done even though Impala does not have TLS/SSL enabled.

Cloudera Bug: OPSAPS-27146

[The Cloudera Manager displays a spurious validation warning about the Cloudera Management Service truststore](#)

Cloudera Manager incorrectly warns that Cloudera Management Service daemons will use HTTPS for communication with either Cloudera Manager or CDH services even if no Cloudera Management Service truststore is in use.

Cloudera Bug: OPSAPS-27473

Issues Fixed in Cloudera Manager 5.3.4

[Slowness observed when accessing replication page in Cloudera Manager](#)

When you access the replication page in Cloudera Manager, the page responds slowly due to a large number of replication history records. The number of displayed historical records has been changed from 100 to 20.

Cloudera Bug: OPSAPS-26529

[Cloudera Manager overwrites the krb5.conf file after disabling management by Cloudera Manager](#)

When a cluster has been configured to enable Kerberos by clicking the **Manage krb5.conf through Cloudera Manager** button, Cloudera Manager writes out a `krb5.conf` file. If a user subsequently disables this feature by disabling the **Manage krb5.conf through Cloudera Manager** option and then restarting Cloudera Manager and the Agent, Cloudera Manager overwrites the existing `krb5.conf` file.

Cloudera Bug: OPSAPS-25282

[Header injection vulnerability in internal logging call](#)

An internal call to a servlet with a malformed logger name created information that could be used in a cross-site scripting attack.

Cloudera Bug: OPSAPS-26487

[Graph for "Total Containers Running Across NodeManagers" displays high values](#)

The graph **Total Containers Running Across NodeManagers**, which displays on the YARN status page, displayed incorrect high values.

Cloudera Bug: OPSAPS-26247

Clicking the "Revert to default" icon stores the default value as a user-defined value in the new configuration pages

Saving an empty configuration value causes the value to be replaced by the default value. The empty value is now saved instead of the default value.

Cloudera Bug: OPSAPS-26340

Some Operational Reports do not return results

The following reports do not return results:

- Overpopulated Directories
- Large Directories
- Custom Reports where the **Replication** parameter is set to 0.

Cloudera Bug: OPSAPS-26178

Xalan has been upgraded to version 2.7.2

To address a vulnerability identified by [CVE-2014-0107](#), Xalan has been updated to version 2.7.2.

Cloudera Bug: OPSAPS-26408

Setting threshold values of -1 or -2 not accepted in new configuration layout pages

When you set threshold values in various properties, you could not select **Specify** and then enter -1 to indicate "Any" or -2 to indicate "Never". You can now enter -1 or -2.

Cloudera Bug: OPSAPS-26066

Rolling upgrade arguments are reversed

The following arguments for rolling upgrade are reversed:

- Sleep seconds
- Failure threshold

Cloudera Bug: OPSAPS-26158

Issues Fixed in Cloudera Manager 5.3.3

hive.metastore.client.socket.timeout default value changed to 60

The default value of the **hive.metastore.client.socket.timeout** property has changed to 60 seconds.

Cloudera Bug: OPSAPS-25270

TLS/SSL Enablement property name changes

The property **hadoop.ssl.enabled** is deprecated. Cloudera Manager has been updated to use either **dfs.http.policy** or **yarn.http.policy** properties instead.

Cloudera Bug: OPSAPS-24985

Changing the Service Monitor Client Config Overrides property requires restart

Cloudera Manager no longer requires you to restart your cluster after changing the **Service Monitor Client Config Overrides** property for a service.

Cloudera Bug: OPSAPS-25272

Cluster name changed from specified name to "cluster" after upgrade

After updating to a new release, Cloudera Manager replaces the specified cluster name with `cluster`. Cloudera Manager now uses the correct cluster name.

Cloudera Manager 5 Release Notes

Cloudera Bug: OPSAPS-25279

Configuration without `host_id` in upgrade DDL causes upgrade problems

A client configuration row in the database DDL did not set `host_id`, causing upgrade problems. Cloudera Manager now catches this condition before upgrading.

Cloudera Bug: OPSAPS-25321

`hive.log.explain.output` property is hidden

The property `hive.log.explain.output` is known to create instability of Cloudera Manager Agents in some specific circumstances, especially when the hive queries generate extremely large EXPLAIN outputs. Therefore, the property has been hidden from the Cloudera Manager configuration screens. You can still configure the property through the use of advanced configuration snippets.

Cloudera Bug: OPSAPS-25852

Slow staleness calculation can lead to ZooKeeper data loss when new servers are added

In Cloudera Manager 5.x, starting new ZooKeeper Servers shortly after adding them can cause ZooKeeper data loss when the number of new servers exceeds the number of old servers.

Cloudera Bug: OPSAPS-25966

Spark and Spark (standalone) services fail to start if you upgrade to CDH 5.2.x parcels from an older CDH package

Spark and Spark standalone services fail to start if you upgrade to CDH 5.2.x parcels from an older CDH package.

Workaround: After upgrading rest of the services, uninstall the old CDH packages, and then start the Spark service.

Cloudera Bug: OPSAPS-24005

Deploy client configuration across cluster after upgrade from Cloudera Manager 4.x to 5.3

Following a 4.x -> 5.3 upgrade, you must deploy client configuration across the entire cluster before deleting any gateway roles, any services, or any hosts. Otherwise the existing 4.x client configurations may be left registered and orphaned on the hosts where they were deployed, requiring you to manually intervene to delete them.

Cloudera Bug: OPSAPS-24279

Oozie health bad when Oozie is HA, cluster is kerberized, and Cloudera Manager and CDH are upgraded

Oozie health will go bad if high availability is enabled in a kerberized cluster with Cloudera Manager 5.0 and CDH 5.0 and Cloudera Manager and CDH are then upgraded to 5.1 or higher.

Cloudera Bug: OPSAPS-23872

Workaround: Disable Oozie HA and then re-enable HA again.

HDFS/Hive replication fails when replicating to target cluster that runs CDH 4.0 and has Kerberos enabled

Cloudera Bug: OPSAPS-25945

Workaround: None.

Issues Fixed in Cloudera Manager 5.3.2

The Review Changes page sometimes hangs

The **Review Changes** page hangs due to the inability to handle the "File missing" scenario.

Cloudera Bug: OPSAPS-24872

High volume of TGT events against AD server with "bad token" messages

A fix has been made to how Kerberos credential caching is handled by management services, resulting in a reduction in the number of Kerberos Ticket Granting Ticket (TGT) requests from the cluster to a KDC. This would have been noticed as "Bad Token" messages being seen in high volume in KDC logging and unnecessarily causing re-authentication by management services.

Cloudera Bug: OPSAPS-24800

Accumulo missing `kinit` when running with Kerberos

Cloudera Manager is unable to run Accumulo when `hostname` command does not return FQDN of hosts.

Cloudera Bug: OPSAPS-24792

HiveServer2 leaks threads when using impersonation

For CDH 5.3 and higher, Cloudera Manager will configure HiveServer2 to use the HDFS cache even when impersonation is on. For earlier CDH, there were bugs with the cache when impersonation was in use, so it is still disabled.

Cloudera Bug: OPSAPS-24751

Deploying client configurations fails if there are dead hosts present in the cluster

If there are hosts in the cluster where the Cloudera Manager agent heartbeat is not working, then deploying client configurations does not work. Starting with Cloudera Manager 5.3.2, such hosts are ignored while deploying client configurations. When the issues with the host are fixed, Cloudera Manager will show those hosts as having stale client configurations, at which point you can redeploy them.

Cloudera Bug: OPSAPS-24748

Health test monitors free space available on the wrong filesystem

The Cloudera Manager Health Test to monitor free space available for the Cloudera Manager Agent's process directory monitors space on the wrong filesystem. It should monitor the `tmpfs` that the Cloudera Manager Agent creates, but instead monitors the Cloudera Manager Agent working directory.

Cloudera Bug: OPSAPS-24733

Starting ZooKeeper Servers from Service or Instance page fails

Stopped ZooKeeper servers cannot be started from the Service or Instance page, but only from the Role page of the server using the `start` action for the role.

Cloudera Bug: OPSAPS-24725

Flume Metrics page does not render agent metrics

Starting in Cloudera Manager 5.3, some or all Flume component data was missing from the Flume Metrics Details page.

Cloudera Bug: OPSAPS-24525

Broken link to help pages on Chart Builder page

The help icon (question mark) on the Chart Builder page returns a 404 error.

Cloudera Bug: OPSAPS-24424

Import MapReduce configurations to YARN now handles NodeManager vcores and memory

Running the wizard to import MapReduce configurations to YARN will now populate `yarn.nodemanager.resource.cpu-vcores` and `yarn.nodemanager.resource.memory-mb` correctly based on equivalent MapReduce configuration.

Cloudera Bug: OPSAPS-21540; KI added Cloudera Manager5.1.0

Issues Fixed in Cloudera Manager 5.3.1

Deploy client configuration across cluster after upgrade from Cloudera Manager 4.x to 5.3

Following a 4.x -> 5.3 upgrade, you must deploy client configuration across the entire cluster before deleting any gateway roles, any services, or any hosts. Otherwise the existing 4.x client configurations may be left registered and orphaned on the hosts where they were deployed, requiring you to manually intervene to delete them.

Cloudera Bug: OPSAPS-24279

Deploy client configuration across cluster after upgrade from Cloudera Manager 4.x to 5.3

Following a 4.x -> 5.3 upgrade, you must deploy client configuration across the entire cluster before deleting any gateway roles, any services, or any hosts. Otherwise the existing 4.x client configurations may be left registered and orphaned on the hosts where they were deployed, requiring you to manually intervene to delete them.

Cloudera Bug: OPSAPS-24279

Oozie health bad when Oozie is HA, cluster is kerberized, and Cloudera Manager and CDH are upgraded

Oozie health will go bad if high availability is enabled in a kerberized cluster with Cloudera Manager 5.0 and CDH 5.0 and Cloudera Manager and CDH are then upgraded to 5.1 or higher.

Cloudera Bug: OPSAPS-2387

Workaround: Disable Oozie HA and then re-enable HA again.

Deploy client configuration no longer fails after 60 seconds

When configuring a gateway role on a host that already contains a role of the same type—for example, an HDFS gateway on a DataNode—the deploy client configuration command no longer fails after 60 seconds.

Cloudera Bug: OPSAPS-24426

service cloudera-scm-server force_start now works

Cloudera Bug: OPSAPS-24489

After deleting services, the Cloudera Manager Server log no longer contains foreign key constraint failure exceptions

Cloudera Bug: OPSAPS-24377

When using Isilon, Cloudera Manager now sets `mapred_submit_replication` correctly

When EMC Isilon storage is used, there is no DataNode, so you cannot set `mapred_submit_replication` to a number smaller than or equal to the number of DataNodes in the network. Cloudera Manager now does the following when setting `mapred_submit_replication`:

- If using HDFS, sets to a minimum of 1 and issues a warning when greater than the number of DataNodes
- If using Isilon, sets to 1 and does not check against the number of DataNodes

Cloudera Bug: OPSAPS-24391

The Cloudera Manager Agent now sets the file descriptor `ulimit` correctly on Ubuntu

Cloudera Bug: OPSAPS-24416

During upgrade, bootstrapping the standby NameNode step no longer fails with standby NameNode connection refused when connecting to active NameNode

Cloudera Bug: OPSAPS-24074

Deploy `krb5.conf` now also deploys it on hosts with Cloudera Management Service roles

Cloudera Bug: OPSAPS-21329

Cloudera Manager allows upgrades to unknown CDH maintenance releases

Cloudera Manager 5.3.0 supports any CDH release less than or equal to 5.3, even if the release did not exist when Cloudera Manager 5.3.0 was released. For packages, you cannot currently use the upgrade wizard to upgrade to such a release. This release adds a custom CDH field for the package case, where you can type in a version that did not exist at the time of the Cloudera Manager release.

Cloudera Bug: OPSAPS-24355

impalad memory limit units error in EnableLlamaRMCommand

The EnableLlamaRMCommand sets the value of the impalad memory limit to equal the NM container memory value. But the latter is in MB, and the former is in bytes. Previously, the command did not perform the conversion; this has been fixed.

Running MapReduce v2 jobs are now visible using the Application Master view

In the Application view, selecting **Application Master** for a MRv2 job previously resulted in no action.

Deleting services no longer results in foreign key constraint exceptions

The Cloudera Manager Server log previously showed several foreign key constraint exceptions that were associated with deleted services. This has been fixed.

HiveServer2 keystore and LDAP group mapping passwords are no longer exposed in client configuration files

The HiveServer2 keystore password and LDAP group mapping passwords were emitted into the client configuration files. This exposed the passwords in plain text in a world-readable file. This has been fixed.

Cloudera Bug: OPSAPS-24442, OPSAPS-24469

A cross-site scripting vulnerability in Cloudera Management Service web UIs fixed

Cloudera Bug: OPSAPS-24080

The high availability wizard now sets the HDFS dependency on ZooKeeper

Cloudera Bug: OPSAPS-24420

Workaround: Before enabling high availability, do the following:

1. Create and start a ZooKeeper service if one does not exist.
2. Go to the HDFS service.
3. Click the **Configuration** tab.
4. Select **HDFS Service-Wide**.
5. Select **Category > Main**.
6. Locate the **ZooKeeper Service** property or search for it by typing its name in the Search box. Select the ZooKeeper service you created.

To apply this configuration property to other role groups as needed, edit the value for the appropriate role group. See [Modifying Configuration Properties Using Cloudera Manager](#).

7. Click **Save Changes** to commit the changes.

BDR no longer assumes superuser is common if clusters have the same realm

If source and destination clusters are in the same Kerberos realm, Cloudera Manager assumed that superuser of the destination is also the superuser on the source cluster. However, HDFS can be configured so that this is not the case.

Cloudera Bug: OPSAPS-24417

Issues Fixed in Cloudera Manager 5.3.0

Setting the default umask in HDFS fails in new configuration layout

Setting the default umask in the HDFS configuration section to 002 in the new configuration layout displays an error: "Could not parse: Default Umask : Could not parse parameter 'dfs_umaskmode'. Was expecting an octal value with a leading 0. Input: 2", preventing the change from being submitted.

Cloudera Bug: OPSAPS-24340

Workaround: Submit the change using the classic configuration layout.

Spark and Spark (standalone) services fail to start if you upgrade to CDH 5.2.x parcels from an older CDH package

Spark and Spark standalone services fail to start if you upgrade to CDH 5.2.x parcels from an older CDH package.

Workaround: After upgrading rest of the services, uninstall the old CDH packages, and then start the Spark service.

Cloudera Bug: OPSAPS-24005

Fixed MapReduce Usage by User reports when using an Oracle database backend

Setting the default umask in HDFS fails in new configuration layout

Setting the default umask in the HDFS configuration section to 002 in the new configuration layout displays an error: "Could not parse: Default Umask : Could not parse parameter 'dfs_umaskmode'. Was expecting an octal value with a leading 0. Input: 2", preventing the change from being submitted.

Workaround: Submit the change using the classic configuration layout.

Cloudera Bug: OPSAPS-24340

Enabling Integrated Resource Management for Impala sets Impala Daemon Memory Limit Incorrectly

The Enable Integrated Resource Management command for Impala (available from the **Actions** pull-down menu on the Impala service page) sets the Impala Daemon Memory Limit to an unusably small value. This can cause Impala queries to fail.

Workaround 1: Upgrade to Cloudera Manager 5.3.

Workaround 2:

1. Run the Enable Integrated Resource Management wizard up to the **Restart Cluster** step. Do not click **Restart Now**.
2. Click on the **leave this wizard** link to exit the wizard without restarting the cluster.
3. Go to the YARN service page. Click **Configuration**, expand the category **NodeManager Default Group**, and click **Resource Management**.
4. Note the value of the Container Memory property.
5. Go to the Impala service page and click **Configuration**. Type `impala daemon memory limit` into the search box.
6. Set the value of the Impala Daemon Memory Limit property to the value noted in step 4 above.
7. Restart the cluster.

Cloudera Bug: OPSAPS-24255

Rolling restart and upgrade of Oozie fails if there is a single Oozie server

Rolling restart and upgrade of Oozie fails if there is only a single Oozie server. Cloudera Manager will show the error message "There is already a pending command on this role."

Workaround: If you have a single Oozie server, do a normal restart.

Cloudera Bug: OPSAPS-23954

Allow "Started but crashed" processes to be restarted by a Start command

In Cloudera Manager 5.3, it is now possible to restart a crashed process with the **Start** command and not just the **Restart** command.

Cloudera Bug: OPSAPS-23567

Add dependency from Agent to Daemons package to yum

In Cloudera Manager 5.3, an explicit dependency has been added from the Agent package to the Daemons package so that upgrading Cloudera Manager 5.2.0 or later to Cloudera Manager 5.3 causes the agent to be upgraded as well. Previously, the Cloudera Manager installer always installed both packages, but this is now enforced at the package dependency level as well.

Cloudera Bug: OPSAPS-24079

Issues Fixed in Cloudera Manager 5.2.8

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera ("Cloudera Products"), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled "Arbitrary remote code execution with InvokerTransformer" and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Issues Fixed in Cloudera Manager 5.2.7

Sensitive Information in Cloudera Manager Diagnostic Support Bundles

Cloudera Manager is designed to transmit certain diagnostic data (or "bundles") to Cloudera. These diagnostic bundles are used by the Cloudera support team to reproduce, debug, and address technical issues for our customers. Cloudera internally discovered a potential vulnerability in this feature, which could cause any sensitive data stored as "advanced configuration snippets (ACS)" (formerly called "safety valves") to be included in diagnostic bundles and transmitted to Cloudera. Notwithstanding any possible transmission, such sensitive data is not used by Cloudera for any purpose.

Cloudera has taken the following actions: (1) modified Cloudera Manager so that it no longer transmits advanced configuration snippets containing the sensitive data, and (2) modified Cloudera Manager TLS/SSL configuration to increase the protection level of the encrypted communication.

Cloudera strives to follow and also help establish best practices for the protection of customer information. In this effort, we continually review and improve our security practices, infrastructure, and data handling policies.

Users affected:

- Users storing sensitive data in advanced configuration snippets

Severity: High

Impact: Possible transmission of sensitive data

CVE: CVE-2015-6495

Immediate Action Required:

- Upgrade Cloudera Manager to one of the following releases: Cloudera Manager 5.4.6, 5.3.7, 5.2.7, 5.1.6, 5.0.7, 4.8.6

Issues Fixed in Cloudera Manager 5.2.6

[Cloudera Manager Agent may become slow or get stuck when responding to commands and when sending heartbeats to Cloudera Manager Server](#)

This issue can occur when Cloudera Navigator auditing is turned on. The auditing code reads audit logs and sends them to the Audit Server. It acquires a lock to protect the list of roles being audited. The same list is also modified by the Cloudera Manager Agent's main thread when a role is started or stopped. If the Audit thread takes too much time to send audits to the Audit Server (which can happen if there is backlog of audit logs), it starves the main Agent thread. This causes the main Agent thread to not send heartbeats and to not respond to commands from the Cloudera Manager Server.

Cloudera Bug: TSB-60

[Problems restarting the NameNode](#)

The NameNode would not restart due to a blocked thread that was processing audit events.

Cloudera Bug: OPSAPS-26129

[Cloudera Manager reports the wrong value for Impala bytes read from cache](#)

Instead of cached bytes, it reported the value of short-circuit bytes.

Cloudera Bug: OPSAPS-26938

[Directory operational reports do not return results](#)

The following reports now return results:

- Overpopulated Directories
- Large Directories
- Custom reports where **Replication= 0**

Cloudera Bug: OPSAPS-26178

[Cloudera Manager uses Xalan 2.7.2](#)

The version of Xalan used in Cloudera Manager has been upgraded to version 2.7.2 to address a possible vulnerability.

Cloudera Bug: CDH-27059

[Cluster name changed to "cluster" after upgrade](#)

After upgrading CDH, the display name of the cluster no longer changes to "cluster".

Cloudera Bug: OPSAPS-25279

[Cloudera Manager monitors the wrong directory for space threshold warnings](#)

Cloudera Manager was monitoring the disk space thresholds using the `/var/run/cloudera-scm-agent` directory. Cloudera Manager now monitors the correct directory: `/var/run/cloudera-scm-agent/process`.

Cloudera Bug: OPSAPS-24529

Default timeout value for the Hive MetaStore changed to 60 seconds

The default value in the Hive **Service Monitor Client Config Overrides** property for the `hive.metastore.client.socket.timeout` property is now 60.

Cloudera Bug: OPSAPS-24346

Changing client configuration overrides

Changes made to client configuration overrides did not take effect until the service was restarted.

Cloudera Bug: OPSAPS-25194

Issues Fixed in Cloudera Manager 5.2.5

Slow staleness calculation can lead to ZooKeeper data loss when new servers are added

In Cloudera Manager 5, starting new ZooKeeper Servers shortly after adding them can cause ZooKeeper data loss when the number of new servers exceeds the number of old servers.

Cloudera Bug: OPSAPS-25966

Permissions set incorrectly on YARN Keytab files

Permissions on YARN Keytab files for NodeManager were set incorrectly to allow read access to any user.

Cloudera Bug: OPSAPS-25390, TSB-48

Issues Fixed in Cloudera Manager 5.2.2

Impalad memory limit units error in EnableLlamaRMCommand has been fixed

The EnableLlamaRMCommand sets the value of the impalad memory limit to equal the NM container memory value. But the latter is in MB, and the former is in bytes. Previously, the command did not perform the conversion; this has been fixed.

Fixed MapReduce Usage by User reports when using an Oracle database backend

HiveServer2 keystore and LDAP group mapping passwords are no longer exposed in client configuration files

The HiveServer2 keystore password and LDAP group mapping passwords were emitted into the client configuration files. This exposed the passwords in plain text in a world-readable file. This has been fixed.

Cloudera Bug: OPSAPS-24442, OPSAPS-24469

Running MapReduce v2 jobs are now visible using the Application Master view

In the Application view, selecting **Application Master** for a MRv2 job previously resulted in no action.

Deleting services no longer results in foreign key constraint exceptions

The Cloudera Manager Server log previously showed several foreign key constraint exceptions that were associated with deleted services. This has been fixed.

Issues Fixed in Cloudera Manager 5.2.1

“POODLE” vulnerability on TLS/SSL enabled ports

The POODLE (Padding Oracle On Downgraded Legacy Encryption) attack takes advantage of a cryptographic flaw in the obsolete TLS/SSLv3 protocol, after first forcing the use of that protocol. The only solution is to disable TLS/SSLv3 entirely. This requires changes across a wide variety of components of CDH and Cloudera Manager in 5.2.0 and all earlier versions. Cloudera Manager 5.2.1 provides these changes for Cloudera Manager 5.2.0 deployments. All Cloudera

Cloudera Manager 5 Release Notes

Manager 5.2.0 users should upgrade to 5.2.1 as soon as possible. For more information, see the [Cloudera Security Bulletin](#).

Can use the log4j advanced configuration snippet to override the default audit logging configuration even if not using Navigator

In Cloudera Manager 5.2.0 only, it was not possible to use the log4j advanced configuration snippet to override the default audit logging configuration when Navigator was not being used.

Cloudera Bug: OPSAPS-23864

Cloudera Manager now collects metrics for CDH 5.0 DataNodes and NameNodes

A number of NameNode and DataNode charts show no data and a number of NameNode and DataNode health checks show unknown results. Metric collection for CDH 5.1 roles is unaffected.

Cloudera Bug: OPSAPS-23363

Workaround: None.

The Reports Manager and Event Server Thrift servers no longer crash on HTTP requests

HTTP queries against the Reports Manager and Event Server Thrift server would earlier cause it to crash with out-of-memory exception.

Cloudera Bug: OPSAPS-23160

Replication commands now use the correct JAVA_HOME if an override has been provided for it

ZooKeeper connection leaks from HBase clients in Service Monitor have been fixed

When a parcel is activated, user home directories are now created with umask 022 instead of using the "useradd" default 077

Cloudera Bug: OPSAPS-23680

Issues Fixed in Cloudera Manager 5.2.0

Bug in openssl-1.0.1e-15 disrupts TLS/SSL communication between Cloudera Manager Agents and CDH services

This issue was observed in TLS/SSL-enabled clusters running CentOS 6.4 and 6.5, where the Cloudera Manager Agent failed when trying to communicate with CDH services. You can see the bug report [here](#).

Cloudera Bug: OPSAPS-23324

Workaround: Upgrade to `openssl-1.0.1e-16.el6_5.7.x86_64`.

Alternatives database points to client configurations of deleted service

In the past, if you created a service, deployed its client configurations, and then deleted that service, the client configurations lived in the alternative database, with a possibly high priority, until cleaned up manually. Now, for a given "alternatives path" (for example `/etc/hadoop/conf`) if there exist both "live" client configurations (ones that would be pushed out with deploy client configurations for active services) and ones that have been "orphaned" client configurations (the service they correspond to has been deleted), the orphaned ones will be removed from the alternatives database. In other words, to trigger cleanup of client configurations associated with a deleted service you must create a service to replace it.

Cloudera Bug: OPSAPS-13336

The YARN property ApplicationMaster Max Retries has no effect in CDH 5

The issue arises because `yarn.resourcemanager.am.max-retries` was replaced with `yarn.resourcemanager.am.max-attempts`.

Cloudera Bug: OPSAPS-21797

Workaround:

1. Add the following to **ResourceManager Advanced Configuration Snippet for yarn-site.xml**, replacing *MAX_ATTEMPTS* with the desired maximum number of attempts:

```
<property>
<name>yarn.resourcemanager.am.max-attempts</name><value>MAX_ATTEMPTS</value>
</property>
```

2. Restart the ResourceManager(s) to pick up the change.

The Spark History Server does not start when Kerberos authentication is enabled.

The Spark History Server does not start when managed by a Cloudera Manager 5.1 instance when Kerberos authentication is enabled.

Workaround:

1. Go to the Spark service.
2. Expand the **Service-Wide > Advanced** category.
3. Add the following configuration to the **History Server Environment Advanced Configuration Snippet** property:

```
SPARK_HISTORY_OPTS=-Dspark.history.kerberos.enabled=true \
-Dspark.history.kerberos.principal=principal \
-Dspark.history.kerberos.keytab=keytab
```

where *principal* is the name of the Kerberos principal to use for the History Server, and *keytab* is the path to the principal's keytab file on the local filesystem of the host running the History Server.

Cloudera Bug: CDH-19867

Hive replication issue with TLS enabled

Hive replication will fail when the source Cloudera Manager instance has TLS enabled, even though the required certificates have been added to the target Cloudera Manager's trust store.

Cloudera Bug: OPSAPS-22159

Workaround: Add the required Certificate Authority or self-signed certificates to the default Java trust store, which is typically a copy of the cacerts file named `jssecacerts` in the `$JAVA_HOME/jre/lib/security/` path of your installed JDK. Use `keytool` to import your private CA certificates into the `jssecacert` file.

The Spark Upload Jar command fails in a secure cluster

The Spark **Upload Jar** command fails in a secure cluster.

Cloudera Bug: OPSAPS-19856

Workaround: To run Spark on YARN, manually upload the Spark assembly jar to HDFS `/user/spark/share/lib`. The Spark assembly jar is located on the local filesystem, typically in `/usr/lib/spark/assembly/lib` or `/opt/cloudera/parcels/CDH/lib/spark/assembly/lib`.

Clients of the JobHistory Server Admin Interface Require Advanced Configuration Snippet

Clients of the JobHistory server administrative interface, such as the `mapred hsadmin` tool, may fail to connect to the server when run on hosts other than the one where the JobHistory server is running.

Cloudera Bug: OPSAPS-20347

Workaround: Add the following to both the **MapReduce Client Advanced Configuration Snippet for mapred-site.xml** and the **Cluster-wide Advanced Configuration Snippet for core-site.xml**, replacing `JOBHISTORY_SERVER_HOST` with the hostname of your JobHistory server:

```
<property>
<name>mapreduce.history.admin.address</name>
<value>JOBHISTORY_SERVER_HOST:10033</value>
</property>
```

Issues Fixed in Cloudera Manager 5.1.7

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera (“Cloudera Products”), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled “Arbitrary remote code execution with InvokerTransformer” and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Issues Fixed in Cloudera Manager 5.1.6

Sensitive Information in Cloudera Manager Diagnostic Support Bundles

Cloudera Manager is designed to transmit certain diagnostic data (or “bundles”) to Cloudera. These diagnostic bundles are used by the Cloudera support team to reproduce, debug, and address technical issues for our customers. Cloudera internally discovered a potential vulnerability in this feature, which could cause any sensitive data stored as “advanced configuration snippets (ACS)” (formerly called “safety valves”) to be included in diagnostic bundles and transmitted to Cloudera. Notwithstanding any possible transmission, such sensitive data is not used by Cloudera for any purpose.

Cloudera has taken the following actions: (1) modified Cloudera Manager so that it no longer transmits advanced configuration snippets containing the sensitive data, and (2) modified Cloudera Manager TLS/SSL configuration to increase the protection level of the encrypted communication.

Cloudera strives to follow and also help establish best practices for the protection of customer information. In this effort, we continually review and improve our security practices, infrastructure, and data handling policies.

Users affected:

- Users storing sensitive data in advanced configuration snippets

Severity: High

Impact: Possible transmission of sensitive data

CVE: CVE-2015-6495

Immediate Action Required:

- Upgrade Cloudera Manager to one of the following releases: Cloudera Manager 5.4.6, 5.3.7, 5.2.7, 5.1.6, 5.0.7, 4.8.6

[Cloudera Manager Agent may become slow or get stuck when responding to commands and when sending heartbeats to Cloudera Manager Server](#)

This issue can occur when Cloudera Navigator auditing is turned on. The auditing code reads audit logs and sends them to the Audit Server. It acquires a lock to protect the list of roles being audited. The same list is also modified by the Cloudera Manager Agent's main thread when a role is started or stopped. If the Audit thread takes too much time to send audits to the Audit Server (which can happen if there is backlog of audit logs), it starves the main Agent thread. This causes the main Agent thread to not send heartbeats and to not respond to commands from the Cloudera Manager Server.

Fixed Issues in Cloudera Manager 5.1.5

[Slow staleness calculation can lead to ZooKeeper data loss when new servers are added](#)

In Cloudera Manager 5, starting new ZooKeeper Servers shortly after adding them can cause ZooKeeper data loss when the number of new servers exceeds the number of old servers.

Cloudera Bug: OPSAPS-25966

[Permissions set incorrectly on YARN Keytab files](#)

Permissions on YARN Keytab files for NodeManager were set incorrectly to allow read access to any user.

Cloudera Bug: OPSAPS-25390, TSB-48

Fixed Issues in Cloudera Manager 5.1.4

["POODLE" vulnerability on TLS/SSL enabled ports](#)

The POODLE (Padding Oracle On Downgraded Legacy Encryption) attack takes advantage of a cryptographic flaw in the obsolete TLS/SSLv3 protocol, after first forcing the use of that protocol. The only solution is to disable TLS/SSLv3 entirely. This requires changes across a wide variety of components of CDH and Cloudera Manager. Cloudera Manager 5.1.4 provides these changes for Cloudera Manager 5.1.x deployments. All Cloudera Manager 5.1.x users should upgrade to 5.1.4 as soon as possible. For more information, see the [Cloudera Security Bulletin](#).

Issues Fixed in Cloudera Manager 5.1.3

[Improved speed and heap usage when deleting hosts on cluster with long history](#)

Speed and heap usage have been improved when deleting hosts on clusters that have been running for a long time.

[When there are multiple clusters, each cluster's topology files and validation for legal topology is limited to hosts in that cluster](#)

When there are multiple clusters, each cluster's topology files and validation for legal topology is limited to hosts in that cluster. Most commands will now fail up front if the cluster's topology is invalid.

[The size of the statement cache has been reduced for Oracle databases](#)

For users of Oracle databases, the size of the statement cache has been reduced to help with memory consumption.

[Improvements to memory usage of "cluster diagnostics collection" for large clusters.](#)

Memory usage of "cluster diagnostics collection" has been improved for large clusters.

Issues Fixed in Cloudera Manager 5.1.2

If a NodeManager that is used as ApplicationMaster is decommissioned, YARN jobs will hang

Cloudera Bug: OPSAPS-21797

Jobs can hang on NodeManager decommission due to a race condition when continuous scheduling is enabled.

Workaround:

1. Go to the YARN service.
2. Expand the **ResourceManager Default Group > Resource Management** category.
3. Uncheck the **Enable Fair Scheduler Continuous Scheduling** checkbox.
4. Click **Save Changes** to commit the changes.
5. Restart the YARN service.

Could not find a healthy host with CDH 5 on it to create HiveServer2 error during upgrade

When upgrading from CDH 4 to CDH 5, if no parcel is active then the error message "Could not find a healthy host with CDH5 on it to create HiveServer2" displays. This can happen when transitioning from packages to parcels, or if you explicitly deactivate the CDH 4 parcel (which is not necessary) before upgrade.

Workaround: Wait 30 seconds and retry the upgrade.

Cloudera Bug: OPSAPS-21942

AWS installation wizard requires Java 7u45 to be installed on Cloudera Manager Server host

Cloudera Manager 5.1 installs Java 7u55 by default. However, the AWS installation wizard does not work with Java 7u55 due to a bug in the jClouds version packaged with Cloudera Manager.

Workaround:

1. Stop the Cloudera Manager Server.

```
sudo service cloudera-scm-server stop
```

2. Uninstall Java 7u55 from the Cloudera Manager Server host.
3. Install Java 7u45 (which you can download from <http://www.oracle.com/technetwork/java/javase/downloads/java-archive-downloads-javase7-521261.html#jdk-7u45-oth-JPR>) on the Cloudera Manager Server host.
4. Start the Cloudera Manager Server.

```
sudo service cloudera-scm-server start
```

5. Run the AWS installation wizard.



Note: Due to a bug in Java 7u45 (http://bugs.java.com/bugdatabase/view_bug.do?bug_id=8014618), TLS/SSL connections between the Cloudera Manager Server and Cloudera Manager Agents and between the Cloudera Management Service and CDH processes break intermittently. If you do not have TLS/SSL enabled on your cluster, there is no impact.

Cloudera Bug: OPSAPS-21855

The YARN property ApplicationMaster Max Retries has no effect in CDH 5

The issue arises because `yarn.resourcemanager.am.max-retries` was replaced with `yarn.resourcemanager.am.max-attempts`.

Cloudera Bug: [OPSAPS-21797](#); KI added Cloudera Manager 5.1.0

Workaround:

1. Add the following to **ResourceManager Advanced Configuration Snippet for yarn-site.xml**, replacing `MAX_ATTEMPTS` with the desired maximum number of attempts:

```
<property>
<name>yarn.resourcemanager.am.max-attempts</name><value>MAX_ATTEMPTS</value>
</property>
```

2. Restart the ResourceManager(s) to pick up the change.

(BDR) Replications can be affected by other replications or commands running at the same time

Replications can be affected by other replications or commands running at the same time, causing replications to fail unexpectedly or even be silently skipped sometimes. When this occurs, a `StaleObjectException` is logged to the Cloudera Manager logs. This is known to occur even with as few as four replications starting at the same time.

Cloudera Bug: OPSAPS-21918

Issues Fixed in Cloudera Manager 5.1.1

Checking "Install Java Unlimited Strength Encryption Policy Files" During Add Cluster or Add/Upgrade Host Wizard on RPM based distributions if JDK 7 or above is pre-installed will cause Cloudera Manager and CDH to fail

If you have manually installed Oracle's official JDK 7 or 8 rpm on a host (or hosts), and check the **Install Java Unlimited Strength Encryption Policy Files** checkbox in the Add Cluster or Add Host wizard when installing Cloudera Manager on that host (or hosts), or when upgrading Cloudera Manager to 5.1, Cloudera Manager installs JDK 6 policy files, which will prevent any Java programs from running against that JDK. Additionally, if this situation does apply, Cloudera Manager/CDH will also choose that particular Java as the default to run against, meaning that Cloudera Manager/CDH fail to start, throwing the following message in logs: `Caused by: java.lang.SecurityException: The jurisdiction policy files are not signed by a trusted signer!`.

Workaround: Do not select the **Install Java Unlimited Strength Encryption Policy Files** checkbox during the aforementioned wizards. Instead download and install them manually, following the instructions on Oracle's website.

- JDK 7 Instructions: <http://www.oracle.com/technetwork/java/javase/downloads/jce-7-download-432124.html>
- JDK 8 Instructions: <http://www.oracle.com/technetwork/java/javase/downloads/jce8-download-2133166.html>



Note:

To return to the default limited strength files, reinstall the original Oracle rpm:

- `yum -y yum reinstall jdk`
- `zypper -zypper in -f jdk`
- `rpm -rpm -iv --replacepkgs filename`, where *filename* is `jdk-7u65-linux-x64.rpm` or `jdk-8u11-linux-x64.rpm`

Cloudera Bug: OPSAPS-21914

Issues Fixed in Cloudera Manager 5.1.0



Important: Cloudera Manager 5.1.0 is no longer available for download from the Cloudera website or from `archive.cloudera.com` due to the JCE policy file issue described in the [Issues Fixed in Cloudera Manager 5.1.1](#) on page 2675 section of the Release Notes. The download URL at `archive.cloudera.com` for Cloudera Manager 5.1.0 now forwards to Cloudera Manager 5.1.1 for the RPM-based distributions for Linux RHEL and SLES.

Changes to property for `yarn.nodemanager.remote-app-log-dir` are not included in the JobHistory Server `yarn-site.xml` and Gateway `yarn-site.xml`

When "Remote App Log Directory" is changed in YARN configuration, the property `yarn.nodemanager.remote-app-log-dir` are not included in the JobHistory Server `yarn-site.xml` and Gateway `yarn-site.xml`.

Cloudera Bug: [OPSAPS-20906](#); KI added Cloudera Manager 5.0.1

Workaround: Set **JobHistory Server Advanced Configuration Snippet (Safety Valve) for yarn-site.xml** and **YARN Client Advanced Configuration Snippet (Safety Valve) for yarn-site.xml** to:

```
<property>
<name>yarn.nodemanager.remote-app-log-dir</name>
<value>/path/to/logs</value>
</property>
```

Secure CDH 4.1 clusters have Hue and Impala share the same Hive

In a secure CDH 4.1 cluster, Hue and Impala cannot share the same Hive instance. If "Bypass Hive Metastore Server" is disabled on the Hive service, then Hue will not be able to talk to Hive. Conversely, if "Bypass Hive Metastore" enabled on the Hive service, then Impala will have a validation error.

Cloudera Bug: OPSAPS-20482

Severity: High

Workaround: Upgrade to CDH 4.2.

The command history has an option to select the number of commands, but does not always return the number you request

Cloudera Bug: OPSAPS-17243

Workaround: None.

Hue does not support YARN ResourceManager High Availability

Cloudera Bug: OPSAPS-18473

Workaround: Configure the Hue Server to point to the active ResourceManager:

1. Go to the Hue service.
2. Click the **Configuration** tab.
3. Select **Scope > Hue or Hue Service-Wide**.
4. Select **Category > Advanced**.
5. Locate the **Hue Server Advanced Configuration Snippet (Safety Valve) for hue_safety_valve_server.ini** property or search for it by typing its name in the Search box.
6. In the **Hue Server Advanced Configuration Snippet for hue_safety_valve_server.ini** field, add the following:

```
[hadoop]
[[ yarn_clusters ]]
[[[default]]]
resourcemanager_host=<hostname of active ResourceManager>
resourcemanager_api_url=http://<hostname of active resource manager>:<web port of active resource manager>
proxy_api_url=http://<hostname of active resource manager>:<web port of active resource manager>
```

The default web port of Resource Manager is 8088.

7. Click **Save Changes** to have these configurations take effect.
8. Restart the Hue service.

Cloudera Manager does not support encrypted shuffle.

Encrypted shuffle has been introduced in CDH 4.1, but it is not currently possible to enable it through Cloudera Manager.

Cloudera Bug: OPSAPS-8480

Severity: Medium

Workaround: None.

[Hive CLI does not work in CDH 4 when "Bypass Hive Metastore Server" is enabled](#)

Hive CLI does not work in CDH 4 when "Bypass Hive Metastore Server" is enabled.

Cloudera Bug: OPSAPS-20721

Workaround: Configure Hive and disable the "Bypass Hive Metastore Server" option.

Alternatively, an approach can be taken that will cause the "Hive Auxiliary JARs Directory" to not work, but will enable basic Hive commands to work. Add the following to "Gateway Client Environment Advanced Configuration Snippet for hive-env.sh," then re-deploy the Hive client configuration:

```
HIVE_AUX_JARS_PATH="
AUX_CLASSPATH=/usr/share/java/mysql-connector-java.jar:/usr/share/java/oracle-connector-java.jar:$(find
/usr/share/cmfd/lib/postgresql-jdbc.jar 2> /dev/null | tail -n 1)
```

[Incorrect Absolute Path to topology.py in Downloaded YARN Client Configuration](#)

The downloaded client configuration for YARN includes the `topology.py` script. The location of this script is given by the `net.topology.script.file.name` property in `core-site.xml`. But the `core-site.xml` file downloaded with the client configuration has an incorrect absolute path to `/etc/hadoop/...` for `topology.py`. This can cause clients that run against this configuration to fail (including Spark clients run in `yarn-client` mode, as well as YARN clients).

Cloudera Bug: OPSAPS-20337

Workaround: Edit `core-site.xml` to change the value of the `net.topology.script.file.name` property to the path where the downloaded copy of `topology.py` is located. This property must be set to an absolute path.

[search_bind_authentication for Hue is not included in .ini file](#)

When `search_bind_authentication` is set to `false`, Cloudera Manager does not include it in `hue.ini`.

Cloudera Bug: OPSAPS-20926

Workaround: Add the following to the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini`:

```
[desktop]
[[ldap]]
search_bind_authentication=false
```

[Erroneous warning displayed on the HBase configuration page on CDH 4.1 in Cloudera Manager 5.0.0](#)

An erroneous "Failed parameter validation" warning is displayed on the HBase configuration page on CDH 4.1 in Cloudera Manager 5.0.0

Severity: Low

Cloudera Bug: OPSAPS-20310

Workaround: Use CDH4.2 or higher, or ignore the warning.

[Host recommissioning and decommissioning should occur independently](#)

In large clusters, when problems appear with a host or role, administrators may choose to decommission the host or role to fix it and then recommission the host or role to put it back in production. Decommissioning, especially host

decommissioning, is slow, hence the importance of parallelization, so that host recommissioning can be initiated before decommissioning is done.

Cloudera Bug: OPSAPS-19055

Fixed Issues in Cloudera Manager 5.0.8

Apache Commons Collections deserialization vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera (“Cloudera Products”), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled “Arbitrary remote code execution with InvokerTransformer” and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, CDH 5.3.8 and lower, CDH 5.2.8 and lower, CDH 5.1.7 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Manager 5.3.8 and lower, and Cloudera Manager 5.2.8 and lower, Cloudera Manager 5.1.6 and lower, Cloudera Manager 5.0.7 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower.

Users affected: All

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Manager 5.5.1 and CDH 5.5.1, Cloudera Manager 5.4.9 and CDH 5.4.9, Cloudera Manager 5.3.9 and CDH 5.3.9, and Cloudera Manager 5.2.9 and CDH 5.2.9, and Cloudera Manager 5.1.7 and CDH 5.1.7, and Cloudera Manager 5.0.8 and CDH 5.0.8.

Fixed Issues in Cloudera Manager 5.0.7

Sensitive Information in Cloudera Manager Diagnostic Support Bundles

Cloudera Manager is designed to transmit certain diagnostic data (or “bundles”) to Cloudera. These diagnostic bundles are used by the Cloudera support team to reproduce, debug, and address technical issues for our customers. Cloudera internally discovered a potential vulnerability in this feature, which could cause any sensitive data stored as “advanced configuration snippets (ACS)” (formerly called “safety valves”) to be included in diagnostic bundles and transmitted to Cloudera. Notwithstanding any possible transmission, such sensitive data is not used by Cloudera for any purpose.

Cloudera has taken the following actions: (1) modified Cloudera Manager so that it no longer transmits advanced configuration snippets containing the sensitive data, and (2) modified Cloudera Manager TLS/SSL configuration to increase the protection level of the encrypted communication.

Cloudera strives to follow and also help establish best practices for the protection of customer information. In this effort, we continually review and improve our security practices, infrastructure, and data handling policies.

Users affected:

- Users storing sensitive data in advanced configuration snippets

Severity: High

Impact: Possible transmission of sensitive data

CVE: CVE-2015-6495

Immediate Action Required:

- Upgrade Cloudera Manager to one of the following releases: Cloudera Manager 5.4.6, 5.3.7, 5.2.7, 5.1.6, 5.0.7, 4.8.6

Cloudera Manager Agent may become slow or get stuck when responding to commands and when sending heartbeats to Cloudera Manager Server

This issue can occur when Cloudera Navigator auditing is turned on. The auditing code reads audit logs and sends them to the Audit Server. It acquires a lock to protect the list of roles being audited. The same list is also modified by the Cloudera Manager Agent's main thread when a role is started or stopped. If the Audit thread takes too much time to send audits to the Audit Server (which can happen if there is backlog of audit logs), it starves the main Agent thread. This causes the main Agent thread to not send heartbeats and to not respond to commands from the Cloudera Manager Server.

Fixed Issues in Cloudera Manager 5.0.6

Slow staleness calculation can lead to ZooKeeper data loss when new servers are added

In Cloudera Manager 5, starting new ZooKeeper Servers shortly after adding them can cause ZooKeeper data loss when the number of new servers exceeds the number of old servers.

Cloudera Bug: OPSAPS-25966

Fixed Issues in Cloudera Manager 5.0.5

“POODLE” vulnerability on TLS/SSL enabled ports

The POODLE (Padding Oracle On Downgraded Legacy Encryption) attack takes advantage of a cryptographic flaw in the obsolete TLS/SSLv3 protocol, after first forcing the use of that protocol. The only solution is to disable TLS/SSLv3 entirely. This requires changes across a wide variety of components of CDH and Cloudera Manager. Cloudera Manager 5.0.5 provides these changes for Cloudera Manager 5.0.x deployments. All Cloudera Manager 5.0.x users should upgrade to 5.0.5 as soon as possible. For more information, see the [Cloudera Security Bulletin](#).

Issues Fixed in Cloudera Manager 5.0.2

Cloudera Manager Impala Query Monitoring does not work with Impala 1.3.1

Impala 1.3.1 contains changes to the runtime profile format that break the Cloudera Manager Query Monitoring feature. This leads to exceptions in the Cloudera Manager Service Monitor logs, and Impala queries no longer appear in the Cloudera Manager UI or API. The issue affects Cloudera Manager 5.0 and 4.6 - 4.8.2.

Cloudera Bug: [OPSAPS-20744](#); KI added Cloudera Manager 5.0.1

Workaround: None. The issue will be fixed in Cloudera Manager 4.8.3 and Cloudera Manager 5.0.1. To avoid the Service Monitor exceptions, turn off the Cloudera Manager Query Monitoring feature by going to **Impala Daemon > Monitoring** and setting the Query Monitoring Period to 0 seconds. Note that the Impala Daemons must be restarted when changing this setting, and the setting must be restored once the fix is deployed to turn the query monitoring feature back on. Impala queries will then appear again in Cloudera Manager's Impala query monitoring feature.

Issues Fixed in Cloudera Manager 5.0.1

Upgrade from Cloudera Manager 5.0.0 beta 1 or beta 2 to Cloudera Manager 5.0.0 requires assistance from Cloudera Support

Contact Cloudera Support before upgrading from Cloudera Manager 5.0.0 beta 1 or beta 2 to Cloudera Manager 5.0.0.

Cloudera Bug: OPSAPS-19894

Workaround: Contact Cloudera Support.

Failure of HDFS Replication between clusters with YARN

HDFS replication between clusters in different Kerberos realms fails when using YARN if the target cluster is CDH 5.

Cloudera Bug: OPSAPS-19930

Workaround: Use MapReduce (MRv1) instead of YARN.

If installing CDH 4 packages, the Impala 1.3.0 option does not work because Impala 1.3 is not yet released for CDH 4.

If installing CDH 4 packages, the Impala 1.3.0 option listed in the install wizard does not work because Impala 1.3.0 is not yet released for CDH 4.

Cloudera Bug: OPSAPS-20406

Workaround: Install using parcels (where the unreleased version of Impala does not appear), or select a different version of Impala when installing with packages.

When updating dynamic resource pools, Cloudera Manager updates roles but may fail to update role information displayed in the UI

When updating dynamic resource pools, Cloudera Manager automatically refreshes the affected roles, but they sometimes get marked incorrectly as running with outdated configurations and requiring a refresh.

Cloudera Bug: OPSAPS-19863

Workaround: Invoke the **Refresh Cluster** command from the cluster actions drop-down menu.

Upgrade of secure cluster requires installation of JCE policy files

When upgrading a secure cluster via Cloudera Manager, the upgrade initially fails due to the JDK not having Java Cryptography Extension (JCE) unlimited strength policy files. This is because Cloudera Manager installs a copy of the Java 7 JDK during the upgrade, which does not include the unlimited strength policy files. To ensure that unlimited strength functionality continues to work, install the unlimited strength JCE policy files immediately after completing the Cloudera Manager Upgrade Wizard and before taking any other actions in Cloudera Manager.

Cloudera Bug: OPSAPS-18641

Workaround: Install the unlimited strength JCE policy files immediately after completing the Cloudera Manager Upgrade Wizard and before taking any other action in Cloudera Manager.

The Details page for MapReduce jobs displays the wrong id for YARN-based replications

The **Details** link for MapReduce jobs is wrong for YARN-based replications.

Cloudera Bug: OPSAPS-19874

Workaround: Find the job id in the link and then go to the **YARN Applications** page and look for the job there.

Reset non-default HDFS File Block Storage Location Timeout value after upgrade from CDH 4 to CDH 5

During an upgrade from CDH 4 to CDH 5, if the HDFS File Block Storage Locations Timeout was previously set to a custom value, it will now be set to 10 seconds or the custom value, whichever is higher. This is required for Impala to start in CDH 5, and any value under 10 seconds is now a validation error. This configuration is only emitted for Impala and no services should be adversely impacted.

Cloudera Bug: OPSAPS-20515

Workaround: None.

HDFS NFS gateway works only on RHEL and similar systems

HDFS NFS gateway works as shipped ("out of the box") only on RHEL-compatible systems, but not on SLES, Ubuntu, or Debian. Because of a bug in native versions of `portmap/rpcbind`, the HDFS NFS gateway does not work out of the box on SLES, Ubuntu, or Debian systems when CDH has been installed from the command-line, using packages. It does work on supported versions of RHEL-compatible systems on which `rpcbind-0.2.0-10.el6` or later is installed, and it does work if you use Cloudera Manager to install CDH, or if you start the gateway as root.

Bug: [731542](#) (Red Hat), [823364](#) (SLES), [594880](#) (Debian)

Workarounds and caveats:

- On Red Hat and similar systems, make sure `rpcbind-0.2.0-10.el6` or later is installed.
- On SLES, Debian, and Ubuntu systems, do one of the following:
 - Install CDH using Cloudera Manager; *or*
 - As of CDH 5.1, start the NFS gateway as root; *or*
 - [Start the NFS gateway without using packages](#); *or*
 - You can use the gateway by running `rpcbind` in insecure mode, using the `-i` option, but keep in mind that this allows anyone from a remote host to bind to the portmap.

Sensitive configuration values exposed in Cloudera Manager

Certain configuration values that are stored in Cloudera Manager are considered sensitive, such as database passwords. These configuration values should be inaccessible to non-administrator users, and this is enforced in the Cloudera Manager Administration Console. However, these configuration values are not redacted when they are read through the API, possibly making them accessible to users who should not have such access.

Cloudera Bug: OPSAPS-20782

Gateway role configurations not respected when deploying client configurations

Gateway configurations set for gateway role groups other than the default one or at the role level were not being respected.

Cloudera Bug: OPSAPS-9853

Documentation reflects requirement to enable at least Level 1 encryption before enabling Kerberos authentication

Cloudera Security documentation now indicates that before enabling Kerberos authentication you should first enable at least Level 1 encryption.

HDFS NFS gateway does not work on all Cloudera-supported platforms

The NFS gateway cannot be started on some Cloudera-supported platforms.

Cloudera Bug: OPSAPS-19957

Workaround: None. Fixed in Cloudera Manager 5.0.1.

Replace `YARN_HOME` with `HADOOP_YARN_HOME` during upgrade

If `yarn.application.classpath` was set to a non-default value on a CDH 4 cluster, and that cluster is upgraded to CDH 5, the classpath is not updated to reflect that `$YARN_HOME` was replaced with `$HADOOP_YARN_HOME`. This will cause YARN jobs to fail.

Cloudera Bug: OPSAPS-20348

Workaround: Reset `yarn.application.classpath` to the default, then re-apply your classpath customizations if needed.

Insufficient password hashing in Cloudera Manager

In versions of Cloudera Manager earlier than 4.8.3 and earlier than 5.0.1, user passwords are only hashed once. Passwords should be hashed multiple times to increase the cost of dictionary based attacks, where an attacker tries many candidate passwords to find a match. The issue only affects user accounts that are stored in the Cloudera Manager database. User accounts that are managed externally (for example, with LDAP or Active Directory) are not affected.

In addition, because of this issue, Cloudera Manager 4.8.3 cannot be upgraded to Cloudera Manager 5.0.0. Cloudera Manager 4.8.3 must be upgraded to 5.0.1 or later.

Cloudera Bug: OPSAPS-20537

Workaround: Upgrade to Cloudera Manager 5.0.1.

Cloudera Manager 5 Release Notes

Upgrade to Cloudera Manager 5.0.0 from SLES older than Service Pack 3 with PostgreSQL older than 8.4 fails

Upgrading to Cloudera Manager 5.0.0 from SUSE Linux Enterprise Server (SLES) older than Service Pack 3 will fail if the embedded PostgreSQL database is in use and the installed version of PostgreSQL is less than 8.4.

Cloudera Bug: OPSAPS-20696

Workaround: Either migrate away from the embedded PostgreSQL database (use MySQL or Oracle) or upgrade PostgreSQL to 8.4 or greater.

MR1 to MR2 import fails on a secure cluster

When running the MR1 to MR2 import on a secure cluster, YARN jobs will fail to find `container-executor.cfg`.

Cloudera Bug: OPSAPS-18657

Workaround: Restart YARN after the import.

After upgrade from CDH 4 to CDH 5, Oozie is missing workflow extension schemas

After an upgrade from CDH 4 to CDH 5, Oozie does not pick up the new workflow extension schemas automatically. User will need to update `oozie.service.SchemaService.wf.ext.schemas` manually and add the schemas added in CDH 5: `shell-action-0.3.xsd`, `sqoop-action-0.4.xsd`, `distcp-action-0.2.xsd`, `oozie-sla-0.1.xsd`, `oozie-sla-0.2.xsd`. Note: None of the existing jobs will be affected by this bug, only new workflows that require new schemas.

Cloudera Bug: OPSAPS-20246

Workaround: Add the new workflow extension schemas to Oozie manually by editing `oozie.service.SchemaService.wf.ext.schemas`.

Issues Fixed in Cloudera Manager 5.0.0

HDFS replication does not work from CDH 5 to CDH 4 with different realms

HDFS replication does not work from CDH 5 to CDH 4 with different realms. This is because authentication fails for services in a non-default realm via the WebHdfs API due to a JDK bug. This has been fixed in JDK6-u34 (b03)) and in JDK7.

Cloudera Bug: OPSAPS-2017

Workaround: Use JDK 7 or upgrade JDK6 to at least version u34.

The Sqoop Upgrade command in Cloudera Manager may report success even when the upgrade fails

Cloudera Bug: CDH-17316

Workaround: Do one of the following:

- 1. Click the Sqoop service and then the **Instances** tab.
- 2. Click the Sqoop server role then the **Commands** tab.
- 3. Click the **stdout** link and scan for the Sqoop Upgrade command.
- In the All Recent Commands page, select the **stdout** link for latest Sqoop Upgrade command.

Verify that the upgrade did not fail.

Cannot restore a snapshot of a deleted HBase table

If you take a snapshot of an HBase table, and then delete that table in HBase, you will not be able to restore the snapshot.

Cloudera Bug: OPSAPS-16881

Severity: Med

Workaround: Use the "Restore As" command to recreate the table in HBase.

Stop dependent HBase services before enabling HDFS Automatic Failover.

When enabling HDFS Automatic Failover, you need to first stop any dependent HBase services. The Automatic Failover configuration workflow restarts both NameNodes, which could cause HBase to become unavailable.

Cloudera Bug: OPSAPS-9645

Severity: Medium

New schema extensions have been introduced for Oozie in CDH 4.1

In CDH 4.1, Oozie introduced new versions for Hive, Sqoop and workflow schema. To use them, you must add the new schema extensions to the Oozie SchemaService Workflow Extension Schemas configuration property in Cloudera Manager.

Cloudera Bug: OPSAPS-10021

Severity: Low

Workaround: In Cloudera Manager, do the following:

1. Go to the CDH 4 **Oozie** service page.
2. Go to the **Configuration** tab, **View and Edit**.
3. Search for "Oozie Schema". This should show the **Oozie SchemaService Workflow Extension Schemas** property.
4. Add the following to the **Oozie SchemaService Workflow Extension Schemas** property:

```
shell-action-0.2.xsd
hive-action-0.3.xsd
sqoop-action-0.3.xsd
```

5. Save these changes.

YARN Resource Scheduler user FairScheduler rather than FIFO.

Cloudera Manager 5.0.0 sets the default YARN Resource Scheduler to FairScheduler. If a cluster was previously running YARN with the FIFO scheduler, it will be changed to FairScheduler next time YARN restarts. The FairScheduler is only supported with CDH4.2.1 and later, and older clusters may hit failures and need to manually change the scheduler to FIFO or CapacityScheduler.

Cloudera Bug: OPSAPS-13335

Severity: Medium

Workaround: For clusters running CDH 4 prior to CDH 4.2.1:

1. Go the YARN service Configuration page
2. Search for "scheduler.class"
3. Click in the Value field and select the schedule you want to use.
4. Save your changes and restart YARN to update your configurations.

Resource Pools Summary is incorrect if time range is too large.

The Resource Pools Summary does not show correct information if the Time Range selector is set to show 6 hours or more.

Cloudera Bug: OPSAPS-16018

Severity: Medium

Workaround: None.

When running the MR1 to MR2 import on a secure cluster, YARN jobs will fail to find `container-executor.cfg`

Cloudera Bug: OPSAPS-18657

Workaround: Restart YARN after the import steps finish. This causes the file to be created under the YARN configuration path, and the jobs now work.

When upgrading to Cloudera Manager 5.0.0, the "Dynamic Resource Pools" page is not accessible

When upgrading to Cloudera Manager 5.0.0, users will not be able to directly access the "Dynamic Resource Pools" page. Instead, they will be presented with a dialog saying that the Fair Scheduler XML Advanced Configuration Snippet is set.

Cloudera Bug: OPSAPS-18768

Workaround:

1. Go to the YARN service.
2. Click the **Configuration** tab.
3. Select **Scope** > **Resource Manager or YARN Service-Wide**.
4. Select **Category** > **Advanced**.
5. Locate the **Fair Scheduler XML Advanced Configuration Snippet** property or search for it by typing its name in the Search box.
6. Copy the value of the **Fair Scheduler XML Advanced Configuration Snippet** into a file.
7. Clear the value of **Fair Scheduler XML Advanced Configuration Snippet**.
8. Recreate the desired Fair Scheduler allocations in the **Dynamic Resource Pools** page, using the saved file for reference.

New Cloudera Enterprise licensing is not reflected in the wizard and license page

Cloudera Bug: OPSAPS-18241

Workaround: None.

The AWS Cloud wizard fails to install Spark due to missing roles

Cloudera Bug: OPSAPS-18937

Workaround: Do one of the following:

- Use the Installation wizard.
- Open a new window, click the Spark service, click on the **Instances** tab, click **Add**, add all required roles to Spark. Once the roles are successfully added, click the **Retry** button in the Installation wizard.

Spark on YARN requires manual configuration

Spark on YARN requires the following manual configuration to work correctly: modify the YARN Application Classpath by adding `/etc/hadoop/conf`, making it the very first entry.

Cloudera Bug: OPSAPS-19788

Workaround: Add `/etc/hadoop/conf` as the first entry in the YARN Application classpath.

Monitoring works with Solr and Sentry only after configuration updates

Cloudera Manager monitoring does not work out of the box with Solr and Sentry on Cloudera Manager 5. The Solr service is in Bad health, and all Solr Servers have a failing "Solr Server API Liveness" health check.

Cloudera Bug: OPSAPS-18152

Severity: Medium

Workaround: Complete the configuration steps below:

1. Create "HTTP" user and group on all machines in the cluster (with `useradd 'HTTP'` on RHEL-type systems).
2. The instructions that follow this step assume there is no existing Solr Sentry policy file in use. In that case, first create the policy file on `/tmp` and then copy it over to the appropriate location in HDFS that Solr Servers check.

If there is already a Solr Sentry policy in use, it must be modified to add the following `[group] / [role]` entries for 'HTTP'. Create a file (for example, `/tmp/cm-authz-solr-sentry-policy.ini`) with the following contents:

```
[groups]
HTTP = HTTP
[roles]
HTTP = collection = admin->action=query
```

3. Copy this file to the location for the "Sentry Global Policy File" for Solr. The associated config name for this location is `sentry.solr.provider.resource`, and you can see the current value by navigating to the **Sentry** sub-category in the **Service Wide** configuration editing workflow in the Cloudera Manager UI. The default value for this entry is `/user/solr/sentry/sentry-provider.ini`. This refers to a path in HDFS.

4. Check if you have entries in HDFS for the parent(s) directory:

```
sudo -u hdfs hadoop fs -ls /user
```

5. You may need to create the appropriate parent directories if they are not present. For example:

```
sudo -u hdfs hadoop fs -mkdir /user/solr/sentry
```

6. After ensuring the parent directory is present, copy the file created in step 2 to this location, as follows:

```
sudo -u hdfs hadoop fs -put /tmp/cm-authz-solr-sentry-policy.ini
/user/solr/sentry/sentry-provider.ini
```

7. Ensure that this file is owned/readable by the solr user (this is what the Solr Server runs as):

```
sudo -u hdfs hadoop fs -chown solr /user/solr/sentry/sentry-provider.ini
```

8. Restart the Solr service. If both Kerberos and Sentry are being enabled for Solr, the MGMT services also need to be restarted. The Solr Server liveness health checks should clear up once SMON has had a chance to contact the servers and retrieve metrics.

Out-of-memory errors may occur when using the Reports Manager

Out-of-memory errors may occur when using the Cloudera Manager Reports Manager.

Cloudera Bug: OPSAPS-19980

Workaround: Set the value of the "Java Heap Size of Reports Manager" property to at least the size of the HDFS filesystem image (`fsimage`) and restart the Reports Manager.

Applying license key using Internet Explorer 9 and Safari fails

Cloudera Manager is designed to work with IE 9 and above and Safari. However the file upload widget used to upload a license currently does not work with IE 9 or Safari. Therefore, installing an enterprise license does not work.

Cloudera Bug: OPSAPS-17223

Workaround: Use another supported browser.

Issues Fixed in Cloudera Manager 5.0.0 Beta 2

The Sqoop Upgrade command in Cloudera Manager may report success even when the upgrade fails

Cloudera Bug: OPSAPS-18924

Workaround: Do one of the following:

- 1. Click the Sqoop service and then the **Instances** tab.
- 2. Click the Sqoop server role then the **Commands** tab.
- 3. Click the **stdout** link and scan for the Sqoop Upgrade command.

- In the All Recent Commands page, select the **stdout** link for latest Sqoop Upgrade command.

Verify that the upgrade did not fail.

The HDFS Canary Test is disabled for secured CDH 5 services.

Due to a bug in Hadoop's handling of multiple RPC clients with distinct configurations within a single process with Kerberos security enabled, Cloudera Manager will disable the HDFS canary test when security is enabled so as to prevent interference with Cloudera Manager's MapReduce monitoring functionality.

Cloudera Bug: OPSAPS-16537

Severity: Medium

Workaround: None

Not all monitoring configurations are migrated from MR1 to MR2.

When MapReduce v1 configurations are imported for use by YARN (MR2), not all of the monitoring configuration values are currently migrated. Users may need to reconfigure custom values for properties such as thresholds.

Cloudera Bug: OPSAPS-16211

Severity: Medium

Workaround: Manually reconfigure any missing property values.

"Access Denied" may appear for some features after adding a license or starting a trial.

After starting a 60-day trial or installing a license for Enterprise Edition, you may see an "access denied" message when attempting to access certain Enterprise Edition-only features such as the Reports Manager. You need to log out of the Admin Console and log back in to access these features.

Cloudera Bug: OPSAPS-16686

Severity: Low

Workaround: Log out of the Admin Console and log in again.

Hue must set impersonation on when using Impala with impersonation.

When using Impala with impersonation, the `impersonation_enabled` flag must be present and configured in the `hue.ini` file. If impersonation is enabled in Impala (in other words, if Impala is using Sentry) then this flag must be set **true**. If Impala is not using impersonation, it should be set **false** (the default).

Cloudera Bug: OPSAPS-16011

Workaround: Set advanced configuration snippet value for `hue.ini` as follows:

1. Go to the **Hue Service Configuration Advanced Configuration Snippet for `hue_safety_valve.ini`** under the Hue service Configuration settings, **Service-Wide > Advanced** category.
2. Add the following, then uncomment the setting and set the value True or False as appropriate:

```
#####  
# Settings to configure Impala  
#####  
  
[impala]  
....  
# Turn on/off impersonation mechanism when talking to Impala  
## impersonation_enabled=False
```

Cloudera Manager Server may fail to start when upgrading using a PostgreSQL database.

If you're upgrading to Cloudera Manager 5.0.0 beta 1 and you're using a PostgreSQL database, the Cloudera Manager Server may fail to start with a message similar to the following:

```
ERROR [main:dbutil.JavaRunner@57] Exception while executing
com.cloudera.cmf.model.migration.MigrateConfigRevisions
java.lang.RuntimeException: java.sql.SQLException: Batch entry <xxx> insert into REVISIONS
(REVISION_ID, OPTIMISTIC_LOCK_VERSION, USER_ID, TIMESTAMP, MESSAGE) values (...)
was aborted. Call getNextException to see the cause.
```

Cloudera Bug: OPSAPS-16971

Workaround: Use `psql` to connect directly to the server's database and issue the following SQL command:

```
alter table REVISIONS alter column MESSAGE type varchar(1048576);
```

After that, your Cloudera Manager server should start up normally.

Issues Fixed in Cloudera Manager 5.0.0 Beta 1

After an upgrade from Cloudera Manager 4.6.3 to 4.7, Impala does not start.

After an upgrade from Cloudera Manager 4.6.3 to 4.7 when Navigator is used, Impala will fail to start because the Audit Log Directory property has not been set by the upgrade procedure.

Cloudera Bug: OPSAPS-15333; added with Cloudera Manager 4.7.0

Severity: Low.

Workaround: Manually set the property to `/var/log/impalad/audit`.

Cloudera Navigator Data Management Release Notes

These release notes provide information about new features, changed features, resolved issues, and known issues (with workarounds when possible) for Cloudera Navigator Data Management component:

New Features and Changes in Cloudera Navigator Data Management

These two sections on this page describe what's new and what has changed in each Cloudera Navigator release:

- [New Features in Cloudera Navigator](#) on page 2688
- [Changed Features in Cloudera Navigator](#) on page 2709

New Features in Cloudera Navigator

The following sections describe what's new in each Cloudera Navigator release.

What's New in Cloudera Navigator 2.15.2

There are no new features in Cloudera Navigator 2.15.2. See [Issues Fixed in Cloudera Navigator 2.15.2](#) on page 2735. See also [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for all known issues to date.

What's New in Cloudera Navigator 2.15.1

In addition to several new features and [changes to existing features](#) listed below, this release also includes several resolved issues. See [Issues Fixed in Cloudera Navigator 2.15.1](#) on page 2736 for details. See also [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for all known issues to date.

- [New default audit extraction filters](#)
- [Improved HDFS metadata extraction performance](#)
- [Improved metadata purge performance](#)
- [New Cloudera Manager health alerts for Navigator Metadata Server](#)
- [New audit events for Fine Grained Privileges and Ownership](#)

New features include:

- **Improvements to audit extraction filters**

This release introduces changes to default HDFS audit filters to better focus the collection of audit events. The following rules were added or updated:

- All HDFS denied access events are accepted. Previously, HDFS denied access events from filtered users and filtered HDFS directories were not recorded.
- HDFS events from the Hive, Spark, and Impala staging directories are discarded. HDFS events from the following job history directories are also discarded:
 - `/user/history/done_intermediate`
 - `/user/spark/applicationHistory`
 - `/user/spark/spark2ApplicationHistory`
- All HDFS delete and rename operations from directories that are not already filtered are accepted.

In addition at the end of the filter list, Cloudera Manager provides a rule that filters events from HDFS `getFileinfo` operations. By default, this filter rule has no affect. Cloudera recommends that you change this rule to **Discard** to stop capturing events for this HDFS operation that indicates access to file *metadata* but does not indicate access to file *data*. This change alone may provide up to 30% reduction in HDFS audit data size.

When upgrading to this release, you may not see the new filters if you have customized your HDFS audit filters previously:

- If HDFS filters have been changed from the default, the new filters will NOT take effect.
- If HDFS filters have not been changed from the default, the new filters will take effect.

To review the recommended filters and determine further optimizations for your system's requirements, see [Recommended Audit Filters](#) and [Reviewing Default Audit Filters](#).

- **Improved HDFS metadata extraction performance**

Previously, Navigator processing for HDFS metadata extraction could take long enough that Navigator could not finish indexing the `fsimage` before the edit log was checkpointed into a new `fsimage`. This caused extraction to get stuck in a loop of always parsing the `fsimage` in bulk extraction mode (slow) rather than moving into incremental extraction from the edit log (fast). This release includes changes that allow Navigator to avoid this "bulk extraction loop" therefore reducing the duration of HDFS extraction by as much as half for the initial extraction. This change improves the performance of subsequent extractions significantly, both by making it much more likely that extraction will be incremental and from additional performance improvements in this fix.

- **Improved metadata purge performance**

This release includes changes that significantly reduce the time required to complete purging for HDFS entities and relations. The purge behavior is changed such that HDFS entities are not included in the purge if they are referenced as an endpoint in a data flow lineage relation.

- **New Cloudera Manager health alerts for Navigator Metadata Server**

This release includes new checks for Navigator Metadata Server health. If the checks fail, they trigger a health alert in Cloudera Manager for the Navigator Metadata Server. The checks include "Solr Element Count Threshold" and "Solr Relation Count Threshold". These checks trigger an alert if documents in Navigator Metadata Server's embedded Solr collection are threatening to exceed the maximum allowed number for either the element core or the relation core. Typically, if either of these alerts trigger, Navigator has encountered a problem that produces more relations than it should. The triggers let you know that this is happening while it can be addressed efficiently.

- **New audit events for Fine Grained Privileges and Ownership**

Navigator Audit Server receives audit events from Sentry, Hive, and Impala that are produced when a user changes ownership for a database, table, or view. For more information, see Sentry's [Object Ownership](#).

What's New in Cloudera Navigator 2.14.2

There are no new features in Cloudera Navigator 2.14.2.

What's New in Cloudera Navigator 2.14.1

There are no new features in Cloudera Navigator 2.14.1.



Note: This version of Cloudera Navigator has a class-compatibility issue that can cause the Navigator Metadata Server to not start; the problem and work-around are described in [Known Issues](#).

What's New in Cloudera Navigator 2.14.0

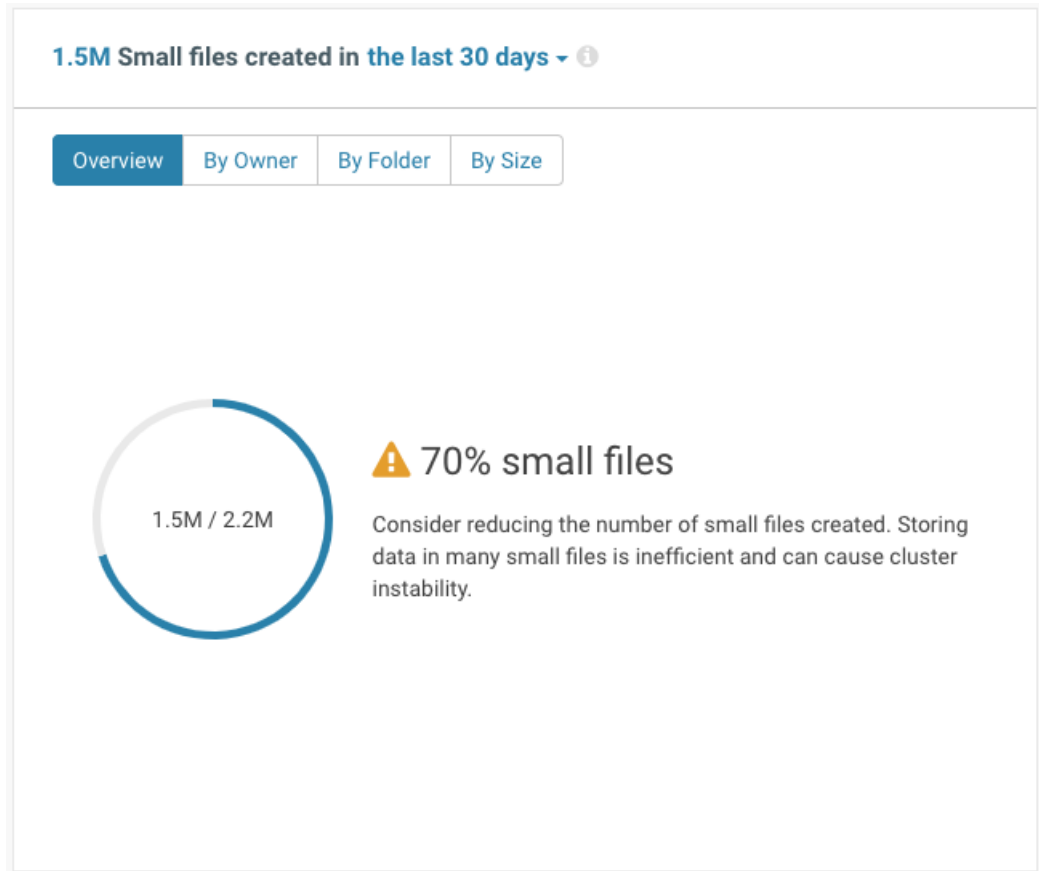
In addition to several new features and [changes to existing features](#) listed below, this release also includes several resolved issues. See [Issues Fixed in Cloudera Navigator 2.14.0](#) on page 2744 for details. See also [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for all known issues to date.

- [Dashboard widget for small files](#)
- [Dashboard widget for active users](#)
- [Preventing concurrent logins](#)

New features include:

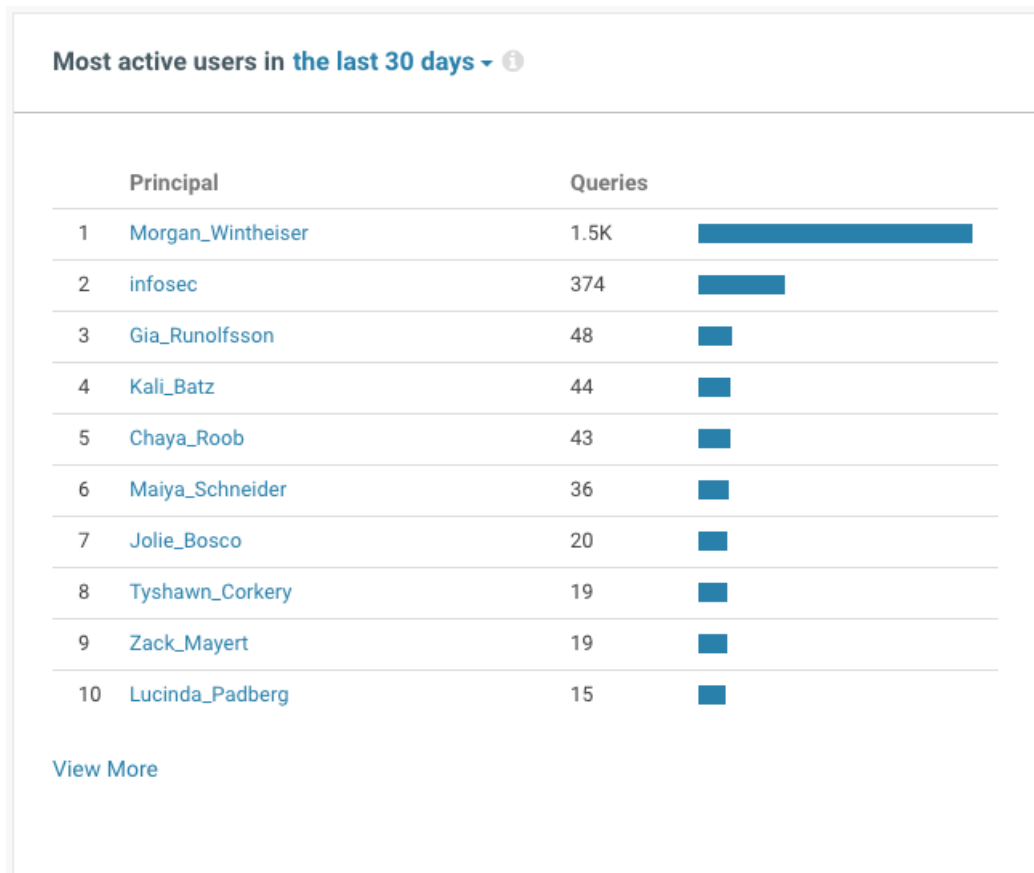
- **Dashboard widget for small files**

Small files can be a performance problem on a Hadoop cluster, particularly when table partitions are assembled from many small files. The Navigator dashboard now provides widgets that visualize the presence of small files based on the files' owner, location, and size. An overview section shows the fraction of small files to large files in the cluster, including a warning if the fraction of small files to larger files exceeds 50%. This metric identifies "small files" as being smaller than 256 KiB.



- **Dashboard widget for active users**

This release includes an additional Navigator dashboard widget that shows the active users by query. The metric includes all Hive and Impala query operation executions.



- **Preventing concurrent logins**

An option is now available to limit the number of simultaneous sessions authenticated against the same user name. The property `nav.max.concurrent.sessions` takes an integer value as a limit; set the value to -1 (default) to turn off the limit. To change the default behavior, in Cloudera Manager, add the property with a new value to the "Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties." Restart the Navigator Metadata Server to apply the change.

What's New in Cloudera Navigator 2.13.4

There are no new features in Cloudera Navigator 2.13.4.

What's New in Cloudera Navigator 2.13.2

In addition to the new features listed below, this release also includes changes and resolved issues. See [What's Changed in Cloudera Navigator 2.13.2](#) on page 2710 for changes, [Issues Fixed in Cloudera Navigator 2.13.2](#) on page 2748 for resolved issues and [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for known issues.

The new feature is:

- **Preventing concurrent logins**

An option is now available to limit the number of simultaneous sessions authenticated against the same user name. The property `nav.max.concurrent.sessions` takes an integer value as a limit; set the value to -1 (default) to turn off the limit. To change the default behavior, in Cloudera Manager, add the property with a new value to the "Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties." Restart the Navigator Metadata Server to apply the change.

Cloudera Navigator Data Management Release Notes

What's New in Cloudera Navigator 2.13.1

This release of Cloudera Navigator includes no new features or fixed issues. See [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for known issues.

What's New in Cloudera Navigator 2.13.0

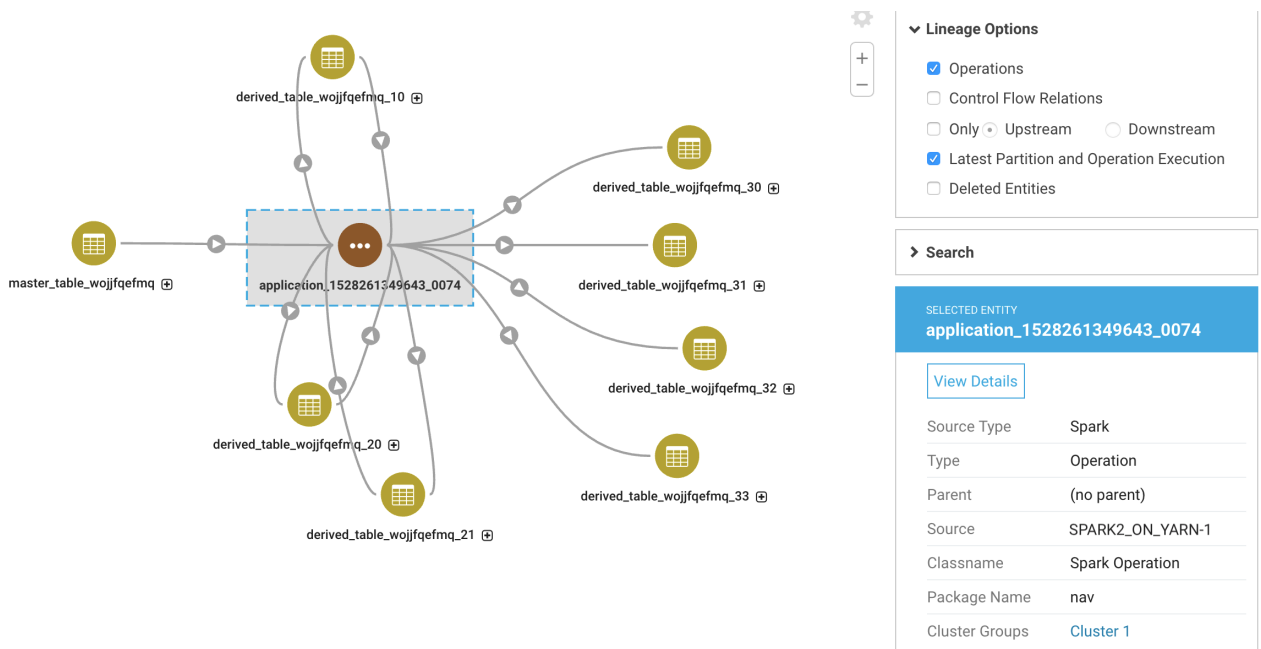
In addition to several new features and [changes to existing features](#) listed below, this release also includes several resolved issues. See [Issues Fixed in Cloudera Navigator 2.13.0](#) on page 2749 for details. See also [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for all known issues to date.

- [Spark2 operations supported in search and lineage](#)
- [Improved Hive extraction performance](#)
- [Group by helps explore connections among search results](#)
- [User interface improvements](#)
- [Audit Event filtering with "Not Like"](#)

New features include:

- **Spark2 operations supported in search and lineage**

This release introduces support for Spark2 operations and operation executions in Navigator metadata management, including search and lineage. Navigator extracts metadata for Spark2 events and generates lineage relations using that information. Spark operations and operation executions appear in Navigator with source type "Spark" and source names that start with "SPARK2_ON_YARN".



All the features of Spark 1 are available for Spark 2, starting with Spark 2 version 2.3.

At this time, lineage for Spark2 is not supported on Altus clusters.

- **Improved Hive extraction performance**

This release includes two areas of performance improvement for Hive metadata extraction:

- **Partitions:** Metadata extraction performance for Hive partitions is considerably improved.
- **Incremental extraction:** Metadata extraction through the Hive Metadata Server (HMS) now only extracts new and changed metadata; the result is more frequent extractions for Hive metadata reducing lag-time for users and significantly less load on HMS.

- **Group by helps explore connections among search results**

Metadata searches in Navigator now include the ability to group search results by common properties. Group by lets you use technical, managed, and custom metadata to quickly identify small files, active SQL users, table-creation trends, and other data aggregation trends revealed by metadata properties.

The screenshot shows the Cloudera Navigator Search interface. On the left, there are filters for SOURCE TYPE (Hive, Pig) and TYPE. The main search results area shows 'Top 7 groups, by "Business Owner" (74 results)'. A 'Group by' dropdown menu is open, showing options like None, Operational, Business Owner (selected), Phase, Retain Until, Sensitivity Management, Custom Metadata, Tags, and Technical Metadata (Compressed, Created, Last Accessed, Last Modified, Owner, Parent Path). The search results table lists groups like RK (19), OX (11), ST (7), AK (3), CF (2), JM (2), and BL (1).

- **User interface improvements**

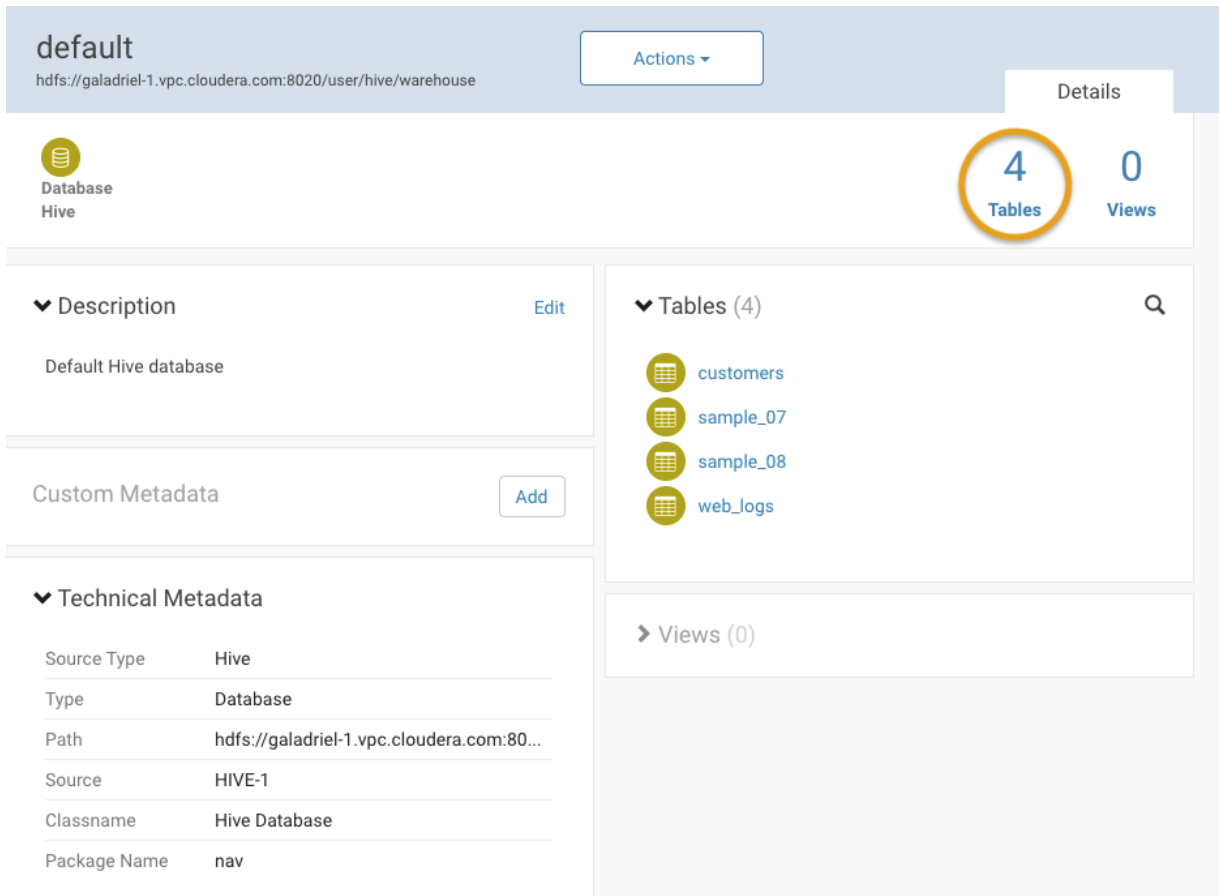
This release includes an improved layout for search results so they are easier to read and understand.

The screenshot shows the Cloudera Navigator Search interface with 372 results. The filters on the left include PERMISSIONS (rwxrwxrwx) and SOURCE TYPE (HDFS). The search results are displayed in a table with columns for Name, Last Modified, Created, Owner, Group, Permissions, Source, and a 'View in Hue' link. The results show HDFS Directory entries for paths like /user/admin with various metadata.

Also, the Navigator console includes some display changes to make detail screens easier to read:

- Tags are now included in the Custom Metadata group.
- Description, Managed Metadata, and Custom Metadata now always show above Technical Metadata, even if the sections are empty.

In addition, on the detail page for a database entity, clicking the table count summary statistic now takes you to search results listing the tables in the database.



- **Audit Event filtering with "Not Like"**

The Audit Events list now allows filtering using the "not like" operator in addition to "like", "equals", and "not equals". Use this operator to remove events from the list based on a partial match such as part of an IP address.

What's New in Cloudera Navigator 2.12.3

In addition to the new feature listed below, this release also includes a change and several resolved issues. See [What's Changed in Cloudera Navigator 2.12.3](#) on page 2710 for changes, [Issues Fixed in Cloudera Navigator 2.12.3](#) on page 2751 for resolved issues and [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for known issues.

The new feature is:

- **Preventing concurrent logins**

An option is now available to limit the number of simultaneous sessions authenticated against the same user name. The property `nav.max.concurrent.sessions` takes an integer value as a limit; set the value to -1 (default) to turn off the limit. To change the default behavior, in Cloudera Manager, add the property with a new value to the "Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties." Restart the Navigator Metadata Server to apply the change.

What's New in Cloudera Navigator 2.12.1

In addition to the new features listed below, this release also includes several resolved issues. See [Issues Fixed in Cloudera Navigator 2.12.1](#) on page 2751 for details. See also [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for all known issues to date.

- [Improved performance for extracting Hive metadata from the HiveServer2 metastore](#)

New features include:

- **Improved performance for extracting Hive metadata from the HiveServer2 metastore.**

This release includes an optimization for extracting metadata for Hive partitions that considerably improves the extraction performance.

What's New in Cloudera Navigator 2.12.0

In addition to several new features and [changes to existing features](#) listed below, this release also includes several resolved issues. See [Issues Fixed in Cloudera Navigator 2.12.0](#) on page 2753 for details. See also [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for all known issues to date.

- [Custom extraction filters for improved performance and reduced storage consumption](#)
- [Support for TLS/SSL for Navigator Audit Server](#)
- [Cloudera Navigator console enhancements and new features](#)

New features include:

- **New custom filters for more selective extraction also help reduce storage consumption.**

Administrators can now create custom blacklists (and whitelists) to specifically exclude (or include) HDFS and Amazon S3 entities during the extraction process. The result is faster extraction—collect only what you really want—and less consumption by Cloudera Navigator on storage media. Users with the Cloudera Manager Full Administrator or Navigator Administrator user roles can configure the settings using the Cloudera Manager Admin Console. The entity filters for HDFS and for Amazon S3 buckets are independently configurable. As shown below, one or more HDFS paths can be blacklisted, while one or more S3 buckets can be whitelisted or blacklisted. The example shows an HDFS path that will be excluded from the extraction process, and two Amazon S3 buckets that will explicitly included in the extraction process (with all other S3 buckets disregarded).

HDFS Filter Enable nav.filter.hdfs.enable	<input checked="" type="checkbox"/> Navigator Metadata Server Default Group ↻
HDFS Filter Blacklist nav.filter.hdfs.blacklist	Navigator Metadata Server Default Group ↻ <input type="text" value="/path-to-exclude"/> <input type="button" value="="/> <input type="button" value="+"/>
S3 Filter Enable nav.filter.s3.enable	<input checked="" type="checkbox"/> Navigator Metadata Server Default Group ↻
S3 Filter Default Action nav.filter.s3.default.action	Navigator Metadata Server Default Group ↻ <input type="radio"/> ACCEPT <input checked="" type="radio"/> DISCARD
S3 Filter list nav.filter.s3.list	Navigator Metadata Server Default Group ↻ <input type="text" value="bucket-one-to-include"/> <input type="button" value="="/> <input type="button" value="+"/> <input type="text" value="bucket-two-to-include"/> <input type="button" value="="/> <input type="button" value="+"/>

Rather than hard-coding paths, regular expressions can be used. For example, to exclude all the files and directories under `/tmp`, include the blacklist filter:

```
/tmp(?:/.*)?
```

Unlike the HDFS filter which functions solely as a blacklist to exclude specified HDFS paths from the extraction process, the S3 Filter list serves as a blacklist or whitelist for the extraction process, depending on the S3 Filter Default Action specified as follows:

- **ACCEPT**—The buckets listed are blacklisted from the extraction process. The extraction process filters out nothing (accepts everything) but ignores the buckets specified in the S3 Filter list.
- **DISCARD**—The extraction process filters out (discards) everything, collecting only from the buckets specified in the S3 Filter list.

Important: When **Discard** is selected as the **S3 Filter Default Action**, be sure to add buckets for extraction to the **S3 Filter list**. Selecting Discard without identifying any buckets in the S3 Filter list results in no data being extracted.

- **TLS/SSL supported by Navigator Audit Server data collection for enhanced security.**

The Navigator Audit Server can now be configured to use TLS/SSL to encrypt communications over HTTP (HTTPS) and from RPC client-server processes (for example, Clouder Manager Agent process used to transmit audit data to the server). As with Navigator Metadata Server, TLS/SSL for Navigator Metadata Server is configured using the Cloudera Manager Admin Console.

- **Cloudera Navigator console new visual element renderings (selectors, icons, fonts, and the like).**

New renderings in the Cloudera Navigator console include the following:

- [New date-and-time selector for Dashboard.](#)
- [New property \(Database\) displays in lineage for tables.](#)
- [New button for SQL copy-and-paste to clipboard for Query Text.](#)
- [New SQL query text formatting for ease of viewing and use.](#)
- [New Feedback button for Google Analytics.](#)

- **New date-and-time selector for Data Stewardship Dashboard.**

The Cloudera Navigator console (**Analytics > Data Stewardship > Dashboard**) now has full date and range picker rather than the drop-down period selector of prior releases. The activity summary number also now shows the sum of events over the given time period (rather than the last point on the graph as was in prior releases), and the time period displays in the graph.

- **Lineage for tables now includes database in Selected Entity.**

In previous releases, the Database, Parent, and other fetched properties displayed on the details page, but not when selected in a lineage graph. In this release, databases are now shown in the Selected Entity box in lineage diagrams when tables are selected.

- **Search results display more quickly in Cloudera Navigator console.**

Counts for selected entities are now being cached for faster Search display. After the initial search for any selected Filter, the facet counts for each facet are populated and the results are cached and refreshed at regular intervals. Thanks to this caching, the Cloudera Navigator console is much more responsive for searching. The larger the cluster, the more noticeable the improvement.

- **Improved formatting of SQL query text for ease of viewing and use.**

The information pop-up for Query Text (on the Details for a selected entity) includes better use of white space and other formatting improvements for enhanced readability. (The screenshot below is based on simple sample code that does not show many of the new and improved visual cues for SQL syntax).

▼ Technical Metadata	
Source Type	Hive
Parent	(no parent)
Query Text	create table sales_by_region as SE...
Inputs	SALESDATA
Outputs	SALES_BY_REGION
Source	HIVE-1
Classname	Hive Query
Package Name	nav
Cluster Groups	Cluster 1

▼ Operation Executions (1)

Query Text 📄

```

create table sales_by_region as
SELECT
  s_neighbor,
  sum(s_price) as Price
FROM
  salesdata
GROUP BY
  s_neighbor

```

➤ Outputs (1)

- **New button to copy SQL to clipboard now available for Query Text.**

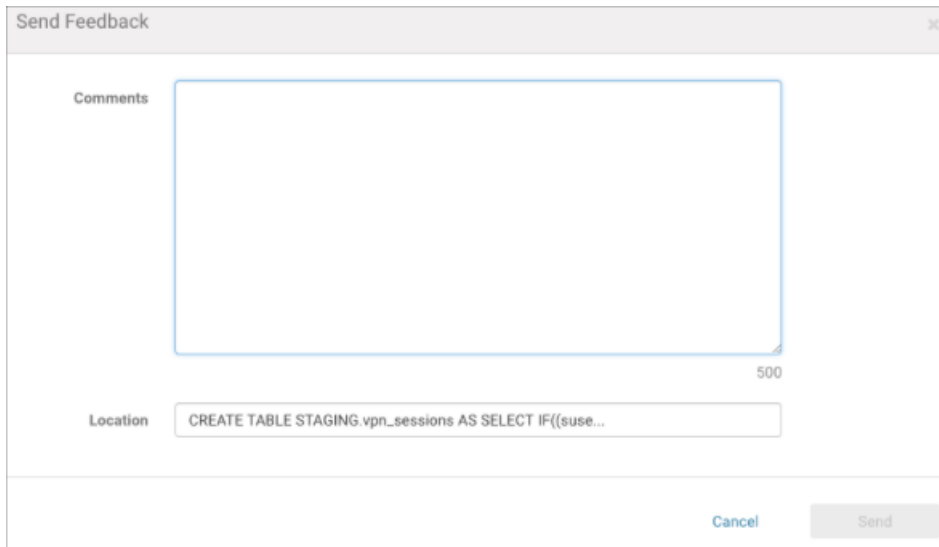
The Query Text pop-up now includes a new copy-to-clipboard button that lets SQL users quickly copy the SQL text displayed in the pop-up to the clipboard for subsequent use elsewhere. As shown in the screenshot above, the location of this button is in the upper right-hand corner of the Query Text pop-up.



Note: Only browsers that support copy-to-clipboard functionality can display this new button. Older browsers that cannot copy-to-clipboard do not display the new button.

- **User feedback module available.**

When Analytics is enabled (Navigator Metadata Server set to **Allow Usage Data Collection**), a **Feedback** button becomes enabled for the Cloudera Navigator console (right-hand side, vertical orientation—see [screenshot above](#)) that lets users log actions and send feedback.



See [What's Changed in Cloudera Navigator 2.12.0](#) on page 2710 for other changes in this release.

What's New in Cloudera Navigator 2.11.2

This release of Cloudera Navigator data management includes no new features. See [Issues Fixed in Cloudera Navigator 2.11.2](#) on page 2754 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for known issues.

What's New in Cloudera Navigator 2.11.1

In addition to support for clusters deployed to Amazon Web Services (AWS) using Cloudera Altus, this release of Cloudera Navigator data management includes some [Issues Fixed in Cloudera Navigator 2.11.1](#) on page 2755 resolved issues. See [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for known issues and workarounds.

- **Obtain metadata and lineage from Altus transient clusters.**

Cloudera Navigator can now extract Hive and MapReduce metadata and lineage from transient clusters deployed to AWS using Cloudera Altus. Entities from clusters instantiated by Cloudera Altus users can be seen in lineage diagrams in a centralized Cloudera Navigator instance running on a persistent or long-running Cloudera Manager cluster.

This new feature is supported with new properties and attributes displayed in the Cloudera Navigator console, to distinguish **long-running** from **Altus** clusters and to distinguish among transient clusters. For example, the **Cluster Group** property identifies all clusters that have been created using the same Altus Environment Name and Altus Cluster name. The **Cluster Instance** distinguishes each instance in the same group from the others, so that metadata and lineage extracted from transient clusters over time can be identified. Here is a summary of new Source Type and Type attributes for clusters:

- Source Type: Cluster
- Cluster-*name* (Cluster Group)
- Deployment Type: Long-running, Altus
- Classname: Cluster Group, Cluster Instance
- Cluster Template, Cluster Instance (Classname)

Entities extracted and their labels as seen in the Cloudera Navigator console include the following:

Entity	Label
altusAwsRegion	Altus AWS Region
altusClusterCrn	Altus Cluster CRN

Entity	Label
altusClusterType	Altus Cluster Type
altusComputeWorkerInstanceType	Altus Compute Worker Instance Type
altusEnvCrn	Altus Env CRN
altusWorkerInstanceType	Altus Worker Instance Type

The **CRN** is the Cloudera Resource Name created and use by Cloudera Altus to manage resources. All nodes in a Cloudera Altus launched cluster have the same CRN.

What's New in Cloudera Navigator 2.11.0

In addition to the new features and the [changes](#) listed below, this release resolves some issues. See [Issues Fixed in Cloudera Navigator 2.11.0](#) on page 2755.

- **Automated Purge feature enabled by default.**

The Purge feature is enabled by default in this release. It uses the default settings shown in the table below unless these settings are changed using the Cloudera Navigator console using the [new Purge Settings](#) tab.



Note: If you enabled the Purge mechanism in a prior release using the Advanced Configuration Snippet (Safety Valve), you should remove it after upgrading to this release. See [What's Changed in Cloudera Navigator 2.11.0](#) on page 2711 for details.

Property	Default	Usage note
How often	Weekly	The Purge process runs automatically each week on the day and time specified.
Day	Saturday	Select a day for the purge that will have minimal impact to your user community.
Time	12 Midnight	Select a time that will have minimal impact on production.
Maximum purge duration	12 hours	Set the amount of time you want to allow for the Purge process to run. The process will not run beyond your specified duration even if it has not completed the purge. Entities purged to that point remain purged. No other Cloudera Navigator operations can occur during the Purge process.
Purge entities deleted more than*	60 days	Select the number of days after entity deletion that will elapse until the purge process removes it. A setting of 60 day purges entities deleted over 61 days ago but retains entities deleted within the last 60 days.
Purge SELECT operations*	Enabled.	Hive and Impala SELECT operations older than the number of days specified in the Only Purge SELECT operations older than* setting are purged.
Only Purge SELECT operations older than*	60 days	The purge will include only those SELECT operations older than this specified number of days.

- **Scalability enhancement for Hive and Impala query execution.**

Heavy Hive and Impala usage can consume lots of storage for similar queries in the Cloudera Navigator data directory (Solr repository). Specifically, because each query can produce (query * query(parts)) entities, a large number of complex queries can result in thousands (or millions) of entities, impeding the lineage tracing process.

To reduce the number of entities created, this release of Cloudera Navigator includes an improved internal query compiler that takes advantage of the fact that many queries vary only by the value of literals. The compiler de-duplicates query executions by dropping the literals when possible and using anonymous queries (using ? rather than literal) instead.



Note: Literals are anonymized in metadata and lineage but remain as literals in audit events subject to your specified audit log redaction policy.

- **Enhancements to Cloudera Navigator console to support new or improved features:**

- **New Administration page for Purge Settings.** Purge Settings is a new tab in the Cloudera Navigator console that lets you schedule weekly Purge operations to remove deleted and stale entities from the Navigator Metadata Server. The Purge is enabled and runs automatically using the [default schedule settings listed in the table](#). Change as needed:

1. Log in to the Cloudera Navigator console.
2. Select **Administration > Purge Settings**.

See [Managing Metadata Storage with Purge](#) in the [Cloudera Navigator Data Management](#) guide for details.

- **Data Stewardship Dashboard.** The Data Stewardship Dashboard now displays:

- Counts of S3 and HDFS objects
- Counts of encrypted and unencrypted objects
- Correct timezone (the server's timezone) despite any differences between timezone of client system displaying the Dashboard and the server

- **Managed Metadata allows deleting namespaces.** The Managed Metadata page now provides a **Manage Namespaces** button that lists namespaces and lets you delete them from the system.

1. Log in to the Cloudera Navigator console.
2. Select **Administration > Managed Metadata**.
3. Click the **Manage Namespaces** button.
4. Click an empty namespace from the list and select **Delete** from the Actions selector.

- **Metadata and lineage extraction from Isilon OneFS.** Dell EMC Isilon OneFS and Cloudera Navigator are now fully integrated, enabling Cloudera Navigator to extract metadata and lineage from clusters that use this advanced storage technology.
- **New configuration setting available for multiple Hue instances behind external load balancers.**

Cloudera Navigator has supported Hue integration (since Cloudera Navigator 2.1 release) by providing links to the appropriate Hue tool in Search results, which open the [Hue interface to the entity](#).

However, for multiple Hue instances deployed behind external load-balancers such as in certain [Hue high availability configurations](#), the links are properly generated only if the host name of the external load-balancer is set as the preferred URL, by configuring the Advanced Configuration Property (Safety Valve) as detailed in the following steps:

1. Log in to Cloudera Manager Admin Console using an account with Full Administrator privileges.
2. Select **Clusters > Cloudera Management Service**.
3. Click the **Configuration** tab.
4. Under filter Scope, click **Navigator Metadata Server**.
5. Under the Category filter, click **Advanced**.
6. Scroll to the property **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**.

- Set the `nav.hue.preferred_baseurl.{clusterName}` to the fully-qualified domain name of your Hue load-balancer using this format:

```
nav.hue.preferred_baseurl.{clusterName}=hue-load-balancer.example.com
```

For example, if a load-balancer for multiple Hue services is setup on host `hue_lb_nginx.example.com`, the entry would be as follows:

Navigator Metadata Server
Advanced Configuration
Snippet (Safety Valve) for
cloudera-
navigator.properties

Navigator Metadata Server Default Group

```
nav.hue.preferred_baseurl.myCluster=https://hue_lb.example.com:8889
```

- Click **Save**.

- Restart the Navigator Metadata Server role. Links to Hue will now be properly created and available in the Search results lists in the Cloudera Navigator console.

- Hue download operations (EXPORT, DOWNLOAD) included in audit data.** Download and export operations from Hue are now collected by Navigator Audit Server. The Audit Event includes the following details:

```
{ "username": "username", "impersonator": "hue", "eventTime": epoch-time-int,
  "operationText": "User username exported to HDFS destination: path/here/filename",
  "service": "service-name", "url": "/service-name-result-path", "allowed": true,
  "operation": "EXPORT", "ipAddress": client-ip-address }
{
```

What's New in Cloudera Navigator 2.10.2

This release of Cloudera Navigator data management includes no new features. See [Issues Fixed in Cloudera Navigator 2.10.2](#) on page 2757 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for known issues.

What's New in Cloudera Navigator 2.10.1

This release of Cloudera Navigator data management includes no new features. See [Issues Fixed in Cloudera Navigator 2.10.1](#) on page 2758 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for known issues.

What's New in Cloudera Navigator 2.10.0

In addition to these new features and [the changes listed below](#), several issues have been resolved in this release. See [Issues Fixed in Cloudera Navigator 2.10.0](#) on page 2759.

- Data Stewardship Dashboard.** The Data Stewardship Dashboard is fully supported and enabled by default in this release. Originally introduced in Cloudera Navigator 2.9 but disabled by default, the Data Stewardship Dashboard displays summary activity about data entities—tables created, tables altered, databases created, SQL queries started, for example. The time frame for the summary is selectable using the Dashboard's trend widget (daily, weekly, monthly, quarterly) to instantly adjust the overall picture of your data. Displayed below the summary results are more detailed views of the data, ranked by count and listing details for the entities as captured by the Navigator Metadata Server. For example, top files by size lists the top 10 files by file name, path, file size, with graphical bar chart rendering for easy visual comparison.
- Data Explorer.** The Data Explorer is an interactive companion to the Data Stewardship Dashboard for viewing and comparing selected data sources over time in terms of averages or trends. Users can select any of the various data sources (databases created, tables created, tables altered, and so on), select chart tools (average, trendline), and select a time period over which to render the data.



Note: To use the Data Stewardship Dashboard and the Data Explorer, select **Analytics > Data Stewardship** from the Cloudera Navigator menu. See [Data Stewardship Dashboard](#) for details.

- **Spark (Spark 1.6) Lineage Support.** The collection mechanism for Spark job lineage has been enhanced to capture Spark application execution—input data, output data, and column-level when possible for each invocation of spark-submit. In addition, lineage for Spark jobs is now collected automatically. Support for this capability in the prior release required manual configuration using a safety valve (advanced configuration snippet). Currently, support is for Spark 1.6 (not Spark 2) and lineage is collected for data read and written using Spark Dataframes and SparkSQL only (RDDs not yet supported). See [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for other limitations.
- **Integration with Navigator for Hue for metadata discovery.** Search and tag partitions, databases, views, tables, columns. This capability must be enabled. See [How to Enable and Use Navigator in Hue](#) for details.
- **Improved logging for Amazon S3 data.** Navigator Server logs now capture information (INFO) about Amazon S3 buckets from which the server is both collecting and not collecting metadata. This enhancement can facilitate better root cause analysis for issues that may arise.

What's New in Cloudera Navigator 2.9.2

This release has no new features. However, some issues have been resolved. See [Issues Fixed in Cloudera Navigator 2.9.2](#) on page 2760 for details. See [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for current listing of all known issues and workarounds.

What's New in Cloudera Navigator 2.9.1

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.9.1](#) on page 2761.

What's New in Cloudera Navigator 2.9.0

- Navigator can now extract technical metadata from data stored on Amazon S3. You can filter results based on the S3 source type and other related S3 properties. See [Navigator and S3](#).
- Solr indexing has been optimized to improve search speed and reduce memory requirements. When Cloudera Navigator is started after an upgrade, you will see a message indicating upgrade progress. See [Upgrading the Cloudera Navigator Management Component](#).
- Managed metadata properties editing has been enhanced. Properties can now be individually viewed, updated, or deleted. Deleted properties are still available from the Administration page and can be restored until they are purged and permanently removed.
- Policy editor has been enhanced:
 - All types of managed metadata can be assigned (text, number, Boolean, date, and enumeration) instead of only text.
 - UI support for managed metadata assignment has been added.
 - Managed metadata assignments can now be multi-valued.
- Auditing health check - You can configure an auditing pipeline health check to verify that auditing is working and is not silently down. See [Verifying that Auditing Is Running](#).
- You can now specify which columns to export when you export an audit report.
- The overall look and feel of the user interface for Cloudera Navigator has been updated.
- Navigator Data Stewardship dashboard - Cloudera Navigator 2.9 introduces the Data Stewardship dashboard, which provides "at-a-glance" information and metrics to help you understand the state of the data and data usage. The dashboard displays information about table, file, and database activity, operations and operation executions, and other information captured by the Navigator Metadata Server. The dashboard is currently unsupported. See [Data Stewardship Dashboard](#).

What's New in Cloudera Navigator 2.8.3

This release of Cloudera Navigator data management includes no new features. See [Issues Fixed in Cloudera Navigator 2.8.3](#) on page 2762 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) on page 2714 for known issues.

What's New in Cloudera Navigator 2.8.2

There are no new features or fixed issues in this release.

What's New in Cloudera Navigator 2.8.1

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.8.1](#) on page 2763.

What's New in Cloudera Navigator 2.8.0

- Support for purging metadata for Hive and Impala SELECT queries, and for YARN, Sqoop, and Pig operations has been added. See [Purging Metadata for HDFS Entities, Hive and Impala Select Queries, and YARN, Sqoop, and Pig Operations](#).
- A new user role that allows editing custom metadata—**Custom Metadata Administrator**—has been added. Grant this role to end users so they can tag their data sets with custom metadata without giving them permission to edit managed metadata. See [User Roles](#).
- You can configure display of inputs and outputs in the entity Details page. See [Enabling Inputs and Outputs to Display](#).

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.8.0](#) on page 2763.

What's New in Cloudera Navigator 2.7.5

There are no new features or fixed issues in this release.

What's New in Cloudera Navigator 2.7.4

There are no new features or fixed issues in this release.

What's New in Cloudera Navigator 2.7.3

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.7.3](#) on page 2764.

What's New in Cloudera Navigator 2.7.2

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.7.2](#) on page 2764.

What's New in Cloudera Navigator 2.7.1

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.7.1](#) on page 2765.

What's New in Cloudera Navigator 2.7.0

- **Platform enhancements**
 - Added support for assigning managed metadata through policies. Only single-valued managed metadata of type Text is currently supported. See [Using Policies to Automate Metadata Tagging](#).
 - Added support for purging physical operations corresponding to logical operations. See [Performing Actions on Entities](#).
- **Cross Site Request Forgery (CSRF)** protection is enabled by default. If you currently use cookies for authentication in API requests with a non-GET method (for example, the POST method), you must also obtain CSRF tokens for the request, or disable CSRF protection.

You can disable CSRF protection by setting `nav.disable_api_csrf_security=true`.

Cloudera Navigator Data Management Release Notes

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.7.0](#) on page 2765

What's New in Cloudera Navigator 2.6.6

There are no new features or fixed issues.

What's New in Cloudera Navigator 2.6.5

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.6.5](#) on page 2768

What's New in Cloudera Navigator 2.6.4

There are no new features or fixed issues.

What's New in Cloudera Navigator 2.6.3

There are no new features or fixed issues.

What's New in Cloudera Navigator 2.6.2

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.6.2](#) on page 2768.

What's New in Cloudera Navigator 2.6.1

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.6.1](#) on page 2769.

What's New in Cloudera Navigator 2.6.0

- **Platform enhancements**

- Added support for defining new types of business metadata.
- Enhanced entity details with type-specific information and behavior.
- Added support for filtering lineage.
- Added support for purging the metadata store.
- Added support for securing audit messages sent to Kafka.

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.6.0](#) on page 2771.

What's New in Cloudera Navigator 2.5.1



Note: In the Cloudera Navigator 2.5.1 release, there are no new features or fixed issues.

What's New in Cloudera Navigator 2.5.0

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.5.0](#) on page 2772.

What's new in Cloudera Navigator 2.4.6



Note: In the Cloudera Navigator 2.4.6 release, there are no new features or fixed issues.

What's New in Cloudera Navigator 2.4.4

Several issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.4.4](#) on page 2772.

What's New in Cloudera Navigator 2.4.3



Note: In the Cloudera Navigator 2.4.3 release, there are no new features or fixed issues.

What's New in Cloudera Navigator 2.4.2



Note: If you are upgrading from a previous version of Cloudera Navigator, Cloudera recommends that you upgrade to version 2.4.3 or higher.

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.4.2](#) on page 2772.

What's New in Cloudera Navigator 2.4.1

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.4.1](#) on page 2774.

What's New in Cloudera Navigator 2.4.0

- **Platform enhancements**

- New entity details page.
- HDFS metadata and audit analytics.
- Command actions, integrated with policies and search, for archiving and purging files.
- Redesigned policies page for scalability.
- Publishing of audit events to Kafka topics. Kafka audit event publishing does not support authorization and Cloudera does not recommend its use in production.
- Navigator SDK for metadata and lineage augmentation.
- Significant scale improvements. Large lineage graphs render faster and consume less memory.

- **Expanded service coverage**

- Hive metadata: support for extended attributes.
- Hive on Spark metadata and lineage.



Important: Like Spark metadata and lineage, Hive on Spark metadata and lineage is not supported or recommended for production use. By default it is disabled. To try this feature, use it in a test environment until Cloudera resolves currently existing issues and limitations to make it ready for production use.

- Oozie metadata and lineage added support for the `hive2` action, which Cloudera recommends that you use in preference to the `hive` action.
- Hue auditing (CDH 5.5.0 and higher): added login and logout, user and group events.
- Cloudera Manager auditing: login and logout events.
- Navigator Metadata Server auditing: successful and failed login events.
- Added filtering of audit events for Sentry, Solr, Impala, and Navigator Metadata Server.

Also see [What's New in Cloudera Manager 5.5.0](#) on page 2495.

What's New in Cloudera Navigator 2.3.10



Note: In the Cloudera Navigator 2.3.10 release, there are no new features or fixed issues.

Cloudera Navigator Data Management Release Notes

What's New in Cloudera Navigator 2.3.9

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.3.9](#) on page 2776.

What's New in Cloudera Navigator 2.3.8

Several issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.3.8](#) on page 2777.



Note: In the Cloudera Navigator 2.3.5, 2.3.6, and 2.3.7 releases, there are no new features or fixed issues.



Note: Although there is a CDH 5.4.4 release, there is no synchronous Cloudera Navigator 2.3.4 release.

What's New in Cloudera Navigator 2.3.3

Several issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.3.3](#) on page 2777.



Note: Although there is a CDH 5.4.2 release, there is no synchronous Cloudera Navigator 2.3.2 release.

What's New in Cloudera Navigator 2.3.1

- Navigator self audit events have been enhanced with additional information such as the names of audit reports and policies
- Performance and stability improvements

Also, a number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.3.1](#) on page 2778.

What's New in Cloudera Navigator 2.3.0

- **Platform enhancements**
 - Redesigned metadata search provides autocomplete, enhanced filtering, and saving searches.
 - Added support for SAML for single sign-on.
- **Expanded service coverage**
 - Added Impala (CDH 5.4 and higher) metadata and lineage
 - Added Cloudera Search (CDH 5.4 and higher) auditing
 - Added auditing for Navigator Metadata Server activity, such as audit views, metadata searches, and policy editing
 - Added support for inferring the schema of HDFS Avro and Parquet entities
 - Added Spark (CDH 5.4 and higher) metadata and lineage.



Important: Spark metadata and lineage is not supported or recommended for production use. By default it is disabled. To try this feature, use it in a test environment until Cloudera resolves currently existing issues and limitations to make it ready for production use.

What's New in Cloudera Navigator 2.2.10



Note: In the Cloudera Navigator 2.2.10 release, there are no new features or fixed issues.

What's New in Cloudera Navigator 2.2.9

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.2.9](#) on page 2779.



Note: Cloudera Navigator releases 2.2.6, 2.2.7, and 2.2.8 have no new features or fixed issues.



Note: Although CDH has a 5.3.5 release, Cloudera Navigator does not have a release for 2.2.5.

What's New in Cloudera Navigator 2.2.4

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.2.4](#) on page 2779.

What's New in Cloudera Navigator 2.2.3

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.2.3](#) on page 2779.

What's New in Cloudera Navigator 2.2.2

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.2.2](#) on page 2780.

What's New in Cloudera Navigator 2.2.1

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.2.1](#) on page 2780.

What's New in Cloudera Navigator 2.2.0

- **Policies** are generally available and are always enabled. Policy properties now support Java expressions.
- **Search** - Search functionality now includes autocomplete.

What's New in Cloudera Navigator 2.1.6

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.1.6](#) on page 2780.

What's New in Cloudera Navigator 2.1.5

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.1.6](#) on page 2780.

What's New in Cloudera Navigator 2.1.4

An issue has been fixed. See [Issues Fixed in Cloudera Navigator 2.1.4](#) on page 2781.



Note: There was no Cloudera Navigator 2.1.3 release.

What's New in Cloudera Navigator 2.1.2

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.1.2](#) on page 2781.

What's New in Cloudera Navigator 2.1.1

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.1.1](#) on page 2782.

What's New in Cloudera Navigator 2.1.0

- **Auditing Component**
 - New tab in the Cloudera Navigator console (the user interface) featuring saved audit reports. The Cloudera Navigator auditing user interface is no longer available from Cloudera Manager. Instead, auditing is integrated with lineage, discovery, and the policy engine. The Cloudera Navigator console is hosted by the Navigator Metadata Server, which is now required for the auditing component.

Cloudera Navigator Data Management Release Notes

- Sentry auditing now includes Sentry commands issued from Impala.

- **Metadata Component**

- Search results include links to the appropriate Hue browser tool for type:

Type	Hue tool
HDFS directories and files	File Browser
Hive database and tables	Hive Metastore (HMS) Manager
MapReduce, YARN, Pig	Job Browser

- **Policies** - support rules for modifying metadata and sending notifications when entities are extracted.



Note: Policies is a beta feature that is disabled by default.

- **Security**

- **Role-Based Access Control** - support assigning groups to roles that constrain access to Navigator features
- **Authentication** - LDAP and Active Directory authentication of Navigator users
- **TLS/SSL** - enable TLS/SSL for encrypted communication

- **API**

- Version changed to v3
- Supports auditing and policies

What's New in Cloudera Navigator 2.0.5

An issue was fixed. See [Issues Fixed in Cloudera Navigator 2.0.5](#) on page 2782.



Note: There is no Cloudera Navigator 2.0.4 release.

What's New in Cloudera Navigator 2.0.3

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.0.3](#) on page 2782.

What's New in Cloudera Navigator 2.0.2

An issue was fixed. See [Issues Fixed in Cloudera Navigator 2.0.2](#) on page 2782.

What's New in Cloudera Navigator 2.0.1

- Masking of personally identifiable information (PII) in query strings that appear in audit events and lineage. Enabled by default.
- REST API support for registering custom metadata for entities before they appear in Navigator.

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator 2.0.1](#) on page 2782.

What's New in Cloudera Navigator 2.0.0

- **Auditing Component**

- Added support for auditing the Sentry service
- Added support for publishing audit logs to syslog

- **Metadata Component**

- Newly designed Query Builder with faceted filtering
 - Simplified Pig lineage
 - Added support for Sqoop and Oozie lineage
 - Many performance and stability improvements
- **Security** - includes Navigator Encrypt and Navigator Key Trustee, formerly known as Gazzang zNcrypt and Gazzang zTrustee. These features provide enterprise-grade encryption and key management. For information on these features, see the [Cloudera Security Datasheet](#) and contact your account team.

Changed Features in Cloudera Navigator

The sections below describe changes to functionality, if any, for each Cloudera Navigator release.

What's Changed in Cloudera Navigator 2.15.2

This maintenance release of Cloudera Navigator data management includes no changed features. See [Issues Fixed in Cloudera Navigator 2.15.2](#) on page 2735 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) for known issues.

What's Changed in Cloudera Navigator 2.15.1

- New Default HDFS audit filters

This release changes the default behavior of HDFS audit event collection ONLY for new installations or upgraded installation if the HDFS audit event filters had not been changed from original default settings.

For more information, see [New Default HDFS audit filters](#) in the What's New section.

- HDFS Analytics disabled by default

The HDFS Analytics available in the Navigator console is now disabled by default. To enable the HDFS Analytics, set the following properties in the Cloudera Management Service configuration:

1. Set `navigator.analytics.enabled` to be true in the **Navigator Audit Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**.
2. Set `nav.analytics.audit.enabled` to be true in the **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**.

What's Changed in Cloudera Navigator 2.14.2

This maintenance release of Cloudera Navigator data management includes no changed features. See [Issues Fixed in Cloudera Navigator 2.14.2](#) on page 2740 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) for known issues.

What's Changed in Cloudera Navigator 2.14.1

This maintenance release of Cloudera Navigator data management includes no changed features. See [Issues Fixed in Cloudera Navigator 2.14.1](#) on page 2742 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) for known issues.

What's Changed in Cloudera Navigator 2.14.0

- **Name and description lengths increased**

This release increases the number of characters available to store entity names and descriptions:

- Name length was 40 characters and is now 500.
- Description length was 500 and is now not limited.

The Navigator console limits manual entry of these values to 40 and 500 characters, so if you edit existing metadata for an entity, you won't be able to save changes if the text is longer than the console limits.

Cloudera Navigator Data Management Release Notes

What's Changed in Cloudera Navigator 2.13.2

- **Name and description lengths increased**

This release increases the number of characters available to store entity names and descriptions:

- Name length was 40 characters and is now 500.
- Description length was 500 and is now not limited.

The Navigator console limits manual entry of these values to 40 and 500 characters, so if you edit existing metadata for an entity, you won't be able to save changes if the text is longer than the console limits.

What's Changed in Cloudera Navigator 2.13.0

- **Changes to Navigator roles and privileges.**

Along with enhancing the metadata management capabilities of Navigator, the Navigator roles and the privileges they describe have changed. The new role names and privileges are described in [User Roles and Privileges Reference](#). One specific change is that the privilege for editing the name and description metadata for Navigator entities is now part of the Managed and Custom Metadata Editor role. Users with that role or the Full Administrator role can add and update entity names and descriptions in the Navigator console.

- **Changes to managed metadata administration.**

The **Administration Managed Metadata** page now shows warnings when properties exist without being associated with entity types. Also, when you add a regular expressions to validate text-type values for managed metadata, the regular expressions are now tested interactively.

What's Changed in Cloudera Navigator 2.12.3

- **Name and description lengths increased**

This release increases the number of characters available to store entity names and descriptions when created or updated through the API:

- Name length was 40 characters and is now 500.
- Description length was 500 and is now not limited.

The Navigator console limits manual entry of these values to 40 and 500 characters, so if you edit existing metadata for an entity, you won't be able to save changes if the text is longer than the console limits.

What's Changed in Cloudera Navigator 2.12.0

- As of Cloudera Manager 5.13, there is a significant change in the volume of HDFS audits that Navigator collects. Cloudera Manager Backup and Disaster Recovery (BDR) feature includes HDFS replication jobs. Previously the jobs used "run as" user to return the access control list (ACL) from the source files; as of 5.13, it uses the source clusters' HDFS principal to get the source file ACL. The change improves the BDR performance but also means that audits that were previously collected for a specific user/owner are now collected for the HDFS principal. If the increase of audits causes problems, consider changing the HDFS audit filter to add HDFS principal to bypass the additional audit records.

- **Deleted entities no longer display by default on Search.**

The Cloudera Navigator console displayed both deleted and non-deleted entities by default (Search did not have a filter for Deleted entities). The default non-filtered Search now shows only non-deleted entities prior to any filtering and a new facet lets user filter out Deleted entities from results or find only Deleted entities.

▼ DELETED	
<input type="radio"/> Show all	1,770
<input checked="" type="radio"/> Do not show Deleted	1,344
<input type="radio"/> Show Deleted only	426

- **Changes to Purge schedule settings.**

Labels and descriptions for the configurable purge options have been clarified. Cloudera Navigator's configurable Purge operation covers operation executions as well as SELECT operations. The labeling and text has been changed to more accurately describe the functionality. The new labels and descriptions are shown below:

Property	Default	Range of selectable values and usage note
Purge HDFS entities deleted more than*	60 days	Select 1 day through 10 days, 20 days through 100 days (in 10-day increments), 150 days, 365 days. These are the number of days after entity deletion that elapse until the purge process removes it. For example, a setting of 1 day purges entities deleted before yesterday but retains entities deleted yesterday.
Purge SELECT operations*	Enabled	Hive and Impala SELECT operations older than days specified in Only Purge SELECT operations older than will be purged.
Purge operations older than*	60 days	Select 10 days through 100 days (10-day increments), 150 days, 365 days. Yarn, Sqoop, and Pig operations older than the specified date will be purged. If Purge SELECT Operations is enabled, Hive and Impala SELECT operations older than the specified date will also be purged.

What's Changed in Cloudera Navigator 2.11.1

This maintenance release of Cloudera Navigator data management includes no changed features. See [Issues Fixed in Cloudera Navigator 2.11.1](#) on page 2755 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) for known issues.

What's Changed in Cloudera Navigator 2.11.0

- **Automated Purge function supplants safety valve setting used in prior releases.** The Administration tab of the Cloudera Navigator console now includes a [Purge Settings](#) page for scheduling the Purge process according to your specifications. If your current system has an Advanced Configuration Snippet (Safety Valve) for `nav.maintenance.purge.schedule` set for a cron job to run purge, you can remove the setting. It will be ignored by the schedule set using the Cloudera Navigator console.

To remove the purge schedule setting from safety valve (if it exists):

1. Log in to Cloudera Manager Admin Console.
2. Select **Clusters > Cloudera Management Service**.
3. Under Scope filter, click **Navigator Metadata Server**.
4. Under Category filter, click **Advanced**.
5. Scroll to the **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**. If a property named `nav.maintenance.purge.schedule` exists and has been set to a cron schedule, remove it. It will be ignored by the schedule set using the Cloudera Navigator console.

What's Changed in Cloudera Navigator 2.10.2

This release of Cloudera Navigator data management includes no changed features. See [Issues Fixed in Cloudera Navigator 2.10.2](#) on page 2757 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) for known issues.

Cloudera Navigator Data Management Release Notes

What's Changed in Cloudera Navigator 2.10.1

This release of Cloudera Navigator data management includes no changed features. See [Issues Fixed in Cloudera Navigator 2.10.1](#) on page 2758 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) for known issues.

What's Changed in Cloudera Navigator 2.10.0

- **Cloudera Navigator Data Stewardship Analytics (Data Stewardship Dashboard) is now fully enabled.** See [Data Stewardship Dashboard](#) (above) for details.
- **Configurable aggregation scheme for metrics.** Every 120 minutes, an internal task runs to extract, transform, and load various counts collected by Navigator to derive metrics and populate the Data Stewardship Dashboard with values. By default, data is aggregated for a 1-day time period.



Note: Prior to Navigator 2.10, the internal ETL task ran every 5 minutes and aggregated 7 days worth of data.

However, the window can be changed by setting the Advanced Configuration Snippet (Safety Valve). The values for new partitions, new tables, and new files depend on the value for this window.

Safety Valve	Description
<code>nav.dashboard.etl.maxDataSetSize=1</code>	Cloudera Manager 5.11/Navigator 2.10 (and later): Default window size is 1 (1 day) Cloudera Manager 5.10/Navigator 2.9 (and prior): Default window size is 7 days
<code>nav.dashboard.etl.include.hdfs=true</code>	Enables the Access Denied field on the Dashboard to display values.

By default, the aggregated Access Denied metric is not populated but you can enable its display by setting the safety valve. To add this safety valve to your system, or to change the default window for the aggregation calculation:

1. Log in to Cloudera Manager Admin Console.
2. Navigate to Cloudera Management Service and then select Navigator Metadata Server (under Scope filter).
3. Click **Configuration** and then click **Service Configuration**.
4. Scroll to the **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**.
5. Enter one or both safety valves into the field as needed. For example:

```
nav.dashboard.etl.maxDataSetSize=3
nav.dashboard.etl.include.hdfs=true
```

What's Changed in Cloudera Navigator 2.9.2

This release of Cloudera Navigator data management includes no changed features. See [Issues Fixed in Cloudera Navigator 2.9.2](#) on page 2760 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) for known issues.

What's Changed in Cloudera Navigator 2.9.1

This release of Cloudera Navigator data management includes no changed features. See [Issues Fixed in Cloudera Navigator 2.9.1](#) on page 2761 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) for known issues.

What's Changed in Cloudera Navigator 2.9.0

Name and description fields, which were previously in custom metadata, have been moved to managed metadata.

What's Changed in Cloudera Navigator 2.8.3

This release of Cloudera Navigator data management includes no changed features. See [Issues Fixed in Cloudera Navigator 2.8.3](#) on page 2762 for resolved issues and see [Known Issues and Workarounds in Cloudera Navigator Data Management](#) for known issues.

What's Changed in Cloudera Navigator 2.7.0

- **Improvement (reduction) to index size by removal of certain types of relations.**

Until now, Cloudera Navigator captured the underlying MapReduce and YARN operations for Hive, Sqoop, and Pig operations, and stored them as physical operations. In this release, previously captured physical operations, operation executions, and associated relations are removed during purge because they were not used in lineage. The upgrade process now starts with Navigator removing the logical-physical relations to operations and operation executions, the relations themselves, their physical endpoints, and any relations connected to those physical endpoints. The end result is improved scalability.

What's Changed in Cloudera Navigator 2.6.0

- Queries generated using filters now use the [Lucene Query Parser](#) , + , - Boolean operator syntax.
- Hive on Spark metadata extraction is now supported.
- The Navigator Metadata Server requires 192 MiB of Java PermGen space instead of 128 MiB. The value of this internal setting used by the JDK is increased automatically when upgrading to Cloudera Manager 5.7 (Cloudera Navigator 2.6).

What's Changed in Cloudera Navigator 2.4.0

- Cloudera Navigator no longer supports JDK 1.6.
- In the Administration tab, after roles have been assigned to at least one group, **Groups with Navigator Roles** is the default selection.

What's Changed in Cloudera Navigator 2.3.3

- In the Search UI, facet values with the count of 0 are not displayed.

What's Changed in Cloudera Navigator 2.3.0

- The PII masking regular expression is superseded by log file redaction in CDH 5.4.

What's Changed in Cloudera Navigator 2.2.0

- **Metadata Component** - Policies created with Cloudera Navigator 2.1 (containing the Beta version policy engine) are not retained when upgrading to Cloudera Navigator 2.2.

What's Changed in Cloudera Navigator 2.1.0

- **Auditing Component**
 - HDFS audit events generated by the `impala` user are discarded by default.

What's Changed in Cloudera Navigator 2.0.1

- **Metadata Component**
 - The REST API version changed to v2.

What's Changed in Cloudera Navigator 2.0.0

- **Auditing Component**

- HDFS audit events generated by the `solr` user are discarded by default.

Known Issues and Workarounds in Cloudera Navigator Data Management

The sections below provide information about current known issues in Cloudera Navigator data management component.

Access, Login, or Cloudera Navigator Console Issues

Navigator Metadata Server fails to start with `InvalidClassException`, local class incompatible

Navigator Metadata Server fails to start (red status in Cloudera Manager) and the Metadata Server log includes the following exception:

```
2018-09-06 10:38:20,786 WARN jndi [main]:
java.io.InvalidClassException: com.mchange.v2.c3p0.WrapperConnectionPoolDataSource;
local class incompatible: stream classdesc serialVersionUID = 7806429541739165290, local
class serialVersionUID = -7086951306718003710
```

Even if Navigator Metadata Server starts once, it may encounter this problem at a later time, depending on which JAR file is loaded.

Workaround: See Knowledge Base Article [Navigator Metadata Server fails to start with "local class incompatible" error](#)

Affects version: Cloudera Manager version 5.15.1 (Navigator 2.14.1)

Fixed version: Cloudera Manager version 5.15.2 (Navigator 2.14.2)


Cloudera Issue: NAV-6724, TSB-342

SAML authentication fails with "Cloudera Manager Only" setting

With the following combination of Cloudera Manager configuration properties set, authentication to Navigator fails:

- Authentication Backend Order: Cloudera Manager Only
- External Authentication Type: SAML

To configure Navigator for SAML authentication, use an Authentication Backend Order other than "Cloudera Manager Only".



Cloudera Issue: NAV-6211

Cloudera Navigator console page display blocked

The Roboto font used in some of the Cloudera Navigator console text is transparently loaded by the console from Google CDN (content delivery network) <https://fonts.googleapis.com/> and <https://fonts.gstatic.com>. If the network environment prevents access to and from these paths, the Cloudera Navigator console cannot load the font, the page hangs, and the console becomes unusable.

Workaround: The workaround depends on the network environment:

- For environments that limit access to and from the public internet, modify the firewall configuration to allow traffic from <https://fonts.googleapis.com/> and <https://fonts.gstatic.com>.

- For environments that restrict access and do not want to allow traffic from the font CDN, modify the firewall configuration to reject access attempts immediately.
- For environments with no DNS access to the public internet, the only workarounds are to obtain the Roboto font and install locally, or hosting your own internal font CDN with fonts from <https://fonts.googleapis.com/> and <https://fonts.gstatic.com> obtained from another network.

Affects version: Cloudera Manager 5.11.1, 5.10.2, 5.10.1, 5.10.0 (Navigator 2.10.1, 2.9.2, 2.9.1, 2.9.0)

Fixed in version: Cloudera Manager 5.12.0, 5.11.2 (Navigator 2.11.0, 2.10.2)

Cloudera Issues: NAV-4871, NAV-4866

Unlimited description through API only

As of Cloudera Manager version 5.15.0 and 5.14.2, name values can be as long as 500 characters and description values are unlimited. However, the Navigator console limits manual entry of the name to 40 characters and descriptions to 500 characters. When you edit the existing metadata for an entity, if the text is longer than the console limits, you won't be able to save changes to the text.

Fixed in version: Cloudera Manager 5.15.0, 5.14.2, 5.13.3 (Navigator 2.14.0, 2.13.2, 2.12.3)

Cloudera Issue: NAV-6228

Internet Explorer 11 document mode issue

When using Microsoft Internet Explorer 11 to access the Cloudera Navigator console, the login page does not display properly due to an issue with that browser's Compatibility View and certain document modes (see the [Microsoft Technet article](#) for details). As of Cloudera Navigator 2.10.1, a warning prompt now displays if the Internet Explorer's document mode is set to IE7 or lower, to alert you to the issue, but prior releases display a blank page rather than the login page.

Workaround: Open the Compatibility View Settings (under the Internet Explorer options menu) and remove the setting for the domain or website. Or use another browser to access the Cloudera Navigator console. Supported browsers include Mozilla Firefox, Google Chrome, or Microsoft Edge.

Affects version: Cloudera Manager 5.11.0, 5.10.x (Navigator 2.10.0, 2.9.x)

Fixed in version: Cloudera Manager 5.11.1 (Navigator 2.10.1)

Cloudera Issues: NAV-3743, NAV-4381

Error message ("not authorized") when logging in

The Cloudera Navigator console saves the state of the last URL accessed when you log out, and opens the same page at your next login. If two different users log in to Navigator using the same browser tab, the state of the first user applies to the second. If the second user does not have permissions to that section of the page, that user receives an error message.

Workaround: Close the browser tab and log in on a new tab. The state is cleared, and the access error message does not appear.

Affects version: Cloudera Manager 5.8.x, 5.7.x (Navigator 2.9.x, 2.8.x)

Fixed in version: Cloudera Manager 5.11 (Navigator 2.10)

Cloudera Issue: NAV-3425

Console doesn't come up, NAV_EXTRACTOR_STATUS doesn't exist

Libraries included in the Navigator release for Cloudera Manager 5.13 were named such that the upgrade marked the schema with the wrong Navigator version. The result is that schema changes between the original version and the final upgrade version are not run. If your Navigator upgrade encounters this problem, the Navigator console does not come up and you'll see the following error in the Navigator Metadata Server log:

```
Table 'nav_meta_user.NAV_EXTRACTOR_STATUS' doesn't exist
```

Cloudera Navigator Data Management Release Notes

Affects version: Cloudera Manager 5.15.0, 5.14.0, 5.13.1, 5.13.0, 5.12.1, 5.12.0 (Navigator 2.14.0, 2.13.0, 2.12.1, 2.12.0, 2.11.1, 2.11.0)

Fixed in version: Cloudera Manager 5.15.1 (Navigator 2.14.1)

Cloudera Issue: NAV-6213

Cloudera Manager Configuration Issues

Adding a blank audit filter removes filter configuration property

In Cloudera Manager, when adding an empty rule to a service's Audit Event Filter and then saving the change, all existing audit event filters are lost. The filter configuration property is removed from Cloudera Manager's list of configuration properties. Reverting the change in the History and Rollback does not restore the previous filters nor reproduce the filter property.

Affected versions: All Cloudera Navigator versions

Fixed in version: N/A

Cloudera Issue: NAV-6096

Overriding safety valve settings disables audit and lineage features

Customers or third party applications such as Unravel may require that `hive.exec.post.hooks` is configured in a HiveServer2 safety valve. Cloudera Manager will comment out the `hive.exec.post.hooks` value that is configured if audit or lineage is enabled for Hive. The safety valve content shows the commented code:

```
<!-- 'hive.exec.post.hooks', originally set to  
'com.cloudera.navigator.audit.hive.HiveExecHookContext,org.apache.hadoop.hive.ql.hooks.LineageLogger'  
(non-final), is overridden below by a safety valve-->
```

This automated change disables Navigator's auditing and lineage features without notification.

Workaround: To fix this problem, manually merge the original HiveServer2 safety valve content for `hive.exec.post.hooks` with the new value. For example, in the case of Unravel, the new safety valve would look like the following:

```
<property>  
  <name>hive.exec.post.hooks</name>  
  
<value>com.unraveldata.flowhive.hook.HivePostHook,com.cloudera.navigator.audit.hive.HiveExecHookContext,org.apache.hadoop.hive.ql.hooks.LineageLogger</value>  
  
  <description>for Unravel, from unraveldata.com</description>  
</property>
```

Cloudera Issue: NAV-5331

Cloudera Manager audit events case-sensitive when using PostgreSQL

When Cloudera Manager and Navigator Audit Server are installed using PostgreSQL databases, the behavior of queries run from the Navigator console is different between the two databases. The result is that Cloudera Manager events are returned only if the query values match the case of the event values as they are stored in the Cloudera Manager database.

For example, use "HiveReplicationCommand" instead of `HIVEREPLICATIONCOMMAND` as it appears in the Audit log.

Cloudera Issue: NAV-6141

Multi-Cluster Environments

Extractors not working in multi-cluster environment

When Navigator Metadata Server is started with a fresh storage directory on versions 5.12.0 or later, extractors are not working when a single Cloudera Manager is managing multiple clusters.

Affects version: Cloudera Manager 5.15.0, 5.14.3, 5.14.2, 5.14.1, 5.14.0, 5.13.1, 5.13.0, 5.12.1, 5.12.0 (Navigator 2.14.0, 2.13.3, 2.13.2, 2.13.1, 2.13.0, 2.12.1, 2.12.0, 2.11.1, 2.11.0)

Fixed in version: Cloudera Manager 5.16.1, 5.15.1, 5.14.4 (Navigator 2.15.1, 2.14.1, 2.13.4)

Cloudera Issue: NAV-6145

Cloudera Navigator is not supported in installations in multi-cluster environments for Cloudera Manager versions 5.12.0, 5.12.1, and 5.13.0

Navigator is not able to uniquely identify more than one cluster in installations of Cloudera Manager that support multiple clusters. This problem only occurs for Cloudera Manager deployments of version 5.12.0, 5.12.1, or 5.13.0. Avoid multi-cluster installations with these releases. This problem does not affect single cluster deployments or Altus clusters.

Affects version: Cloudera Manager 5.13.0, 5.12.1, 5.12.0 (Navigator 2.12.0, 2.11.1, 2.11.0)

Fixed in version: Cloudera Manager 5.14.0, 5.13.1, 5.12.2 (Navigator 2.13.0, 2.12.1, 2.11.2)

Cloudera Issue: NAV-5615

Table-to-HDFS links not established when Navigator supports multiple, high-availability clusters

When Navigator extracts metadata for multiple clusters and when the clusters are configured for high availability operation, Navigator does not correctly link tables to their HDFS backing data. The result is that lineage between Hive tables and their physical data are not created. In addition, some Hive table metadata that is derived from the backing files is not available.

You may see errors in the log such as the following:

```
2018-02-19 09:01:32,999 ERROR com.cloudera.nav.persist.impl.CompositeLinker
[CDHExecutor-0-CDHUrlClassLoader@01010d7e]: Internal error while
linking.java.lang.IllegalArgumentException: expected one element but was:
<com.cloudera.nav.core.model.Source@b22e68ba, com.cloudera.nav.core.model.Source@5b1bac2c>
```

Affects version: Cloudera Manager 5.13.x, 5.12.x (Navigator 2.12.x, 2.11.x)

Fixed in version: Cloudera Manager 5.14.0 (Navigator 2.13.0)

Cloudera Issue: NAV-5749

AWS and Amazon S3

Update AWS Credentials from the same AWS account

After configuring Cloudera Navigator with a specific set of AWS Credentials for Amazon S3, any future changes to the credentials (for example, if you rotate credentials on a regular basis) must be for the same AWS account (IAM user). Changing the AWS Credentials to those of a different IAM user results in errors from the Amazon Simple Queue Service (used transparently by Cloudera Navigator).

Workaround: If a new key is provided to Cloudera Navigator, the key must belong to the same AWS account as the prior key.

Affects version: Cloudera Manager 5.10 and later (Navigator 2.9)

Cloudera Issues: NAV-3990

Unnamed folders on Amazon S3 not extracted

Unnamed folders on Amazon S3 are not extracted by Navigator, but the content of the folders is extracted. For example, a top-level folder the top level folder in the bucket has no name (for example, `/bucket//folder/file`), it is extracted as `/bucket/folder/file`.

Affects version: Cloudera Manager 5.10 and later (Navigator 2.9)

Cloudera Issue: NAV-3981

Cloudera Navigator Data Management Release Notes

Implicit folders are not marked as deleted

If an implicit folder is deleted in Amazon S3, it does not appear as deleted in Cloudera Navigator console.

Workaround: To prevent folders deleted from Amazon S3 from appearing in Navigator Search results, include `implicit:false` in the search query.

Affects version: Cloudera Manager 5.10 and later (Navigator 2.9)

Cloudera Issue: NAV-3802

Extraction delayed while Amazon S3 inconsistent

Inconsistencies that occur in AWS (for example, due to eventual consistency) can delay Navigator extraction of metadata and lineage from Amazon S3. When Cloudera Navigator detects an inconsistency, extraction may stop until the inconsistency is resolved in AWS. Cloudera Navigator will retry at the next scheduled extraction.

Affects version: Cloudera Manager 5.10 and later (Navigator 2.9)

Cloudera Issue: NAV-4028

Microsoft Azure

Less Secure Credentials Protection Policy can expose Azure credentials in audit logs

When you use Cloudera Manager to configure the ADLS Connector service using the Less Secure option for the Credentials Protection Policy, it is possible for Hive audit logs to include Microsoft Azure credentials. If you are using Navigator Audit Server, these credentials may appear in audit reports. To mitigate this problem, make sure that access to Hive logs is appropriately controlled and that Navigator users with Auditing Viewer roles are cleared to have access to the Hive credentials.

Fixed in version: Cloudera Manager 5.16.1, 5.15.1, 5.14.2 (Navigator 2.15.1, 2.14.1, 2.13.2)

Cloudera Issues: NAV-5861, CDH-56241

Data Stewardship Dashboard

Counts displayed in Dashboard and Search may differ

Counts for databases, tables, views, and other entities displayed in Cloudera Navigator Dashboard can differ from Search values for the same objects. The Dashboard displays data from the Navigator Audit Server, which contains actual counts captured continuously. On the other hand, Search (Solr) returns data that has been periodically extracted from HMS (Hive Metadata Server). Tables, databases, views, or other objects created or destroyed between extracts are not reflected in values displayed by Search (Solr), but they are contained in the audit data displayed in the Dashboard.

Fixed in version: Cloudera Manager 5.11.0 (Navigator 2.10.0)

Cloudera Issue: NAV-4192

Tables Populated count reflects INSERT and UPDATE statements

The Tables Populated display widget in the Data Stewardship Dashboard reflects the number of times that a table has been loaded with data, such as through INSERT and UPDATE statements—not the number of unique tables loaded. For example, a single table to which data is added (through 6 INSERT statements) and that has also had 4 UPDATE statements submitted in the same period would report Tables Populated as 10.

Cloudera Issue: NAV-3886

Search and Managed Metadata

Managed metadata property values hidden

Property values are not visible when the property is defined with multivalued enabled and the type was any numeric type.

Affects version: Cloudera Manager 5.14.1, 5.14.0 (Navigator 2.13.1, 2.13.0)

Fixed in version: Cloudera Manager 5.15.0, 5.14.2 (Navigator 2.14.0, 2.13.2)


Cloudera Issue: NAV-5899

Incorrect operation of the Search filter for "Deleted Time"

The Deleted Time search filter does not produce expected results for filtering entities.

Workaround: Specify the delete time range as timestamps in UNIX epoch in milliseconds. For example, to search for entities deleted between January 1, 2017 00:01:00 UTC and January 5, 2017 23:59:00 UTC, enter the following search string in the search box:

```
+deleteTime:[1514764860000 TO 1515196740000]
```



Fixed in version: Cloudera Manager 5.14.0 and 5.13.1 (Navigator 2.13.0 and 2.12.1)

Cloudera Issue: NAV-1168

Hive, Hue, Impala, Pig

Hive lineage missing

Because of race conditions within a single JDBC session, Navigator may not show lineage when Hive queries run in parallel within the same session. Hive maintains lineage per session and it clears all lineage information from the session when a query completes. The result is that any query that does not generate lineage will clear the lineage for any other queries that are running in parallel. It can occur for any JDBC clients, such as Hue. Hue uses same session to run user queries and background refresh queries. If you are using JDBC to connect to Hive and you issue concurrent queries using a single JDBC connection, you may encounter this issue.

Affects version: Cloudera Manager 5.12.0 and later (Navigator 2.14.0)

Fixed in version: Cloudera Manager 6.0.0 (Navigator 6.0.0)

Cloudera Issue: CDH-55693

Hive Operations and Cloudera Navigator Support Matrix

The table below lists Hive DDL and DML statements and whether Cloudera Navigator supports the operation with metadata and lineage extraction.

Hive operations	Cloudera Navigator Support	Comment
Abort	⊘	Operation does not generate data flow.
Alter Table/Partition/Column	✓	Known Issue: NAV-5365. ALTER TABLE RENAME TO does not create query entity. ALTER TABLE CHANGE column name does not create query entity.
Create Table	✓	
Create/Drop Macro	⊘	Operation does not generate data flow.
Create/Drop/Alter Index	⊘	Operation does not generate data flow.

Hive operations	Cloudera Navigator Support	Comment
Create/Drop/Alter View	⊘	Operation does not generate data flow.
Create/Drop/Alter/Use Database	⊘	Operation does not generate data flow.
Create/Drop/Grant/Revoke Roles and Privileges	⊘	Operation does not generate data flow.
Create/Drop/Reload Function	⊘	Operation does not generate data flow.
DELETE	⊘	Requires ACID support. Hive ACID not supported.
Describe	⊘	Operation does not generate data flow.
Drop/Truncate Table	⊘	
EXPORT	⊘	Known Issue. . External tables can be exported to HDFS but Cloudera Navigator does not create a query entity for the EXPORT.
IMPORT	⊘	Known Issue. . External tables can be imported from HDFS but Cloudera Navigator does not create a query entity for the IMPORT.
INSERT data into Hive Tables from queries	✔	
INSERT data into the file system from queries	✔	
INSERT values into tables from SQL	✔	
LOAD	⊘	Known Issue. . LOADING a CSV from HDFS into an existing Hive table does not generate lineage.
MERGE	⊘	Requires ACID support. Hive ACID not supported.
MSCK REPAIR	⊘	Tables track their respective partitions. Queries to create or repair partitions using MSCK are not captured as query entities.
Show	⊘	Operation does not generate data flow.
UPDATE	⊘	Requires ACID support. Hive ACID not supported.

Missing Impala audit events for database-level operations

Impala audit events at the database level are incorrectly recorded, making it seem like they are missing from the Navigator audit data.

Fixed in version: Cloudera Manager 5.14.2, 5.13.2, 5.12.3, and 5.11.3 (Navigator 2.13.2, 2.12.2, 2.11.3, and 2.10.3)

Cloudera Issue: NAV-5648

Impala and Hive audit events fail to be captured when one audit event includes 4-byte characters, such as an emoji

This problem applies when the Navigator Audit Server database is a MySQL database configured to use the "UTF8" character set.

When a query includes an emoji or other Unicode supplementary-plane character that is encoded as four bytes in UTF-8, Navigator Audit Server fails to process the event and any following events from the same service.

Workaround: There are work-arounds for this problem; contact Cloudera Technical Support for details.

Cloudera Issue: NAV-4845

Multi-cluster support for Hive lineage

When there were multiple Hive MetaStore (HMS) sources associated with a single Hive service, the method Navigator used to determine the source was not accurate, resulting in Impala and Spark extractors creating lineage relations that referred to the wrong sources and cross-cluster lineage relations were not generated.

Affects version: Cloudera Manager 5.14.x, 5.13.x, 5.12.x (Navigator 2.13.x, 2.12.x, 2.11.x)

Fixed in version: Cloudera Manager 5.15.0 (Navigator 2.14.0)

Cloudera Issue: NAV-5859

Hive service configuration in auditing component

For Hive services, the auditing component does not support the "Shutdown" option for the "Queue Policy" property.

Severity: Low

Workaround: None.

Cloudera Issue: OPSAPS-11537

Hue service audit log and Unknown IP address

The IP address in a Hue service audit log displays as "unknown".

Severity: Low

Workaround: None.

Cloudera Issue: OPSAPS-11986

Changing auditing configuration for Hive or Hue requires restart

Whenever a change is made to a Hive service's audit configuration, Beeswax must be restarted so that the Hue service audit log can reflect the change.

Severity: Low

Workaround: None.

Cloudera Issue: OPSAPS-12274

Viewing Navigator tags in Hue overloads Metadata Server heap

When viewing Cloudera Navigator tags through Hue, Navigator uses more memory than usual and does not release the memory after logging out of Hue. Eventually, the calls between Hue and Navigator will occupy the majority of the heap space allocated to Navigator Metadata Server.

To work around this problem you may need to restart the Navigator Metadata Server periodically to clear the heap usage.

Affects version: Cloudera Manager 5.11.x and later (Navigator 2.10.x)

Cloudera Issue: NAV-4326

Hue links are missing from Navigator console

The links to Hue from entities in the Navigator console are not generated. Each time the console attempts to show a link, the Navigator Metadata Server log includes an error similar to the following:

```
10:05:32.539 AM ERROR HueUtility
[qtp778450792-57]: Unable to get Hue link from CM.
```

```
java.lang.LinkageError: loader constraint violation: when resolving method
"com.cloudera.nav.server.NavOptions.getPreferredBaseURLForCluster(Ljava/lang/String;)Lcom/google/common/base/Optional;"
the class loader (instance of org/eclipse/jetty/webapp/WebAppClassLoader) of the current
class, com/cloudera/nav/utils/HueUtility, and the class loader (instance of
sun/misc/Launcher$AppClassLoader) for resolved class, com/cloudera/nav/server/NavOptions,
have different Class objects for the type
ter(Ljava/lang/String;)Lcom/google/common/base/Optional; used in the signature{
```

Affects version: Cloudera Manager 5.9.0 and later (Navigator 2.8.0)

Fixed in version: Cloudera Manager 5.13.0 and 5.12.1 (Navigator 2.12.0, 2.11.1)

Severity: Low

Workaround: None.

Cloudera Issue: NAV-4756

Lineage not generated for Pig operations on Hive tables using HCatalog loader

When accessing a Hive table using Pig, lineage is generated in Navigator when using physical file loads, such as:

```
A = LOAD '/user/hive/warehouse/navigator_demo.db/salesdata';
B = LIMIT A 16;
STORE B INTO '/user/hive/warehouse/navigator_demo.db/salesdata_sample_file' using
PigStorage ('');
```

However, when accessing the Hive table using the HCatalog load, lineage for the Pig operation is not generated when browsing the source table lineage. Such as:

```
A = LOAD 'navigator_demo.salesdata' using org.apache.hive.hcatalog.pig.HCatLoader();
B = LIMIT A 16;
STORE B INTO 'navigator_demo.salesdata_sample_hcatalog' using
org.apache.hive.hcatalog.pig.HCatStorer();
```

Cloudera Issue: NAV-3411

HiveServer1 and Hive CLI support deprecated

Cloudera Navigator requires HiveServer2 for complete governance Hive queries. Cloudera Navigator does not capture audit events for queries that are run on HiveServer1/Hive CLI, and lineage is not captured for certain types of operations that are run on HiveServer1.

[HiveServer1 and Hive CLI are deprecated in CDH 5](#). If you use Cloudera Navigator to capture auditing, lineage, and metadata for Hive operations, upgrade to HiveServer2 if you have not done so already.

Cloudera Issue: TSB-185

Navigator Audit Server

Cloudera Manager role name change causes audit events to not be captured

When the role name for a Cloudera Manager service changes, it can cause the Navigator Audit Server to not process the events for that service. You can identify this problem by reviewing the contents of the write-ahead log (*_audit_wal) for the service. This log includes the list of audit log files that are queued to be sent to Navigator. The name of the audit log files includes the service name and the role name; if you see that the structure of the audit log file names have changed at some point, it may be that the name change is preventing the audit server from processing the audits. You can find the *_audit_wal file among the log files for the service. For example, the Hive Server2 *_audit_wal file is typically in the following location on the node where the service is running:

```
/var/log/hive/audit/hiveserver2*_audit_wal
```

Workaround: To restart audit event processing for the service:

1. Remove the *_audit_wal file for the service.

2. Move all audit log files that include the older role name out of the `audit` directory.
3. [Restart the Cloudera Manager agent](#) for that service.

The Navigator audit plugin will recreate the `*audit_wal` file from the remaining audit log files and Navigator Audit Server will collect the audit events.

Cloudera Issues: NAV-6126, NAV-6205

Logging Threshold setting is not honored

The value for Navigator Audit Server Logging Threshold found in Cloudera Manager is not honored. Instead, messages are logged at trace level and displayed at DEBUG syslog level. This configuration property is set in **Cloudera Management Service > Configuration > Navigator Audit Server**.

Affects version: Cloudera Manager 5.12.x and later (Navigator 2.11.x)

Cloudera Issue: NAV-3737

Audit reports download option shows incorrect API version

The API version is not correct in the Cloudera Navigator download option for audit reports.

Cloudera Issue: NAV-3699

Publishing to Kafka fails after upgrade

This issue affects only those Cloudera Navigator instances that publish events to Kafka, and only when upgrading from Cloudera Manager 5.11 to Cloudera Manager 5.11.1.

After upgrading to a new release of Cloudera Manager, Cloudera Navigator Audit Server fails to publish to Kafka. The log (`cloudera-scm-server.log`) contains errors such as the following:

```
kafka.network.InvalidRequestException: Error getting request for apiKey: 3 and apiVersion:
 2
at kafka.network.RequestChannel$Request.liftedTree2$1(RequestChannel.scala:95)
at kafka.network.RequestChannel$Request.<init>(RequestChannel.scala:87)
at
at kafka.network.Processor$$anonfun$processCompletedReceives$1.apply(SocketServer.scala:488)
at
at kafka.network.Processor$$anonfun$processCompletedReceives$1.apply(SocketServer.scala:483)
at scala.collection.Iterator$class.foreach(Iterator.scala:893)
at scala.collection.AbstractIterator.foreach(Iterator.scala:1336)
. . .
```

The issue is due to upgraded Kafka client jar files that conflict with the logging (log4j) infrastructure.

Workaround:

1. Open a browser and navigate to the Cloudera Kafka client and log4j repositories:
 - [Kafka clients](#)
 - [Kafka log4j appenders](#)
2. Download the version of the Kafka jar files that is compatible with your Cloudera Navigator Audit Server. For example, for Cloudera Manager 5.11.0/Cloudera Navigator 2.10.0, download:
 - [kafka-clients-0.9.0-kafka-2.0.0.jar](#) (634.16 KB)
 - [kafka-log4j-appender-0.9.0-kafka-2.0.0.jar](#) (8.47 KB)
3. On the host system running the Navigator Audit Server, remove the symbolic links to the Kafka client jar files located in the `/usr/share/cmf/cloudera-navigator-audit-server` directory. The symlinks are defined as follows:
 - `kafka-log4j-appender-0.10.1.0.jar -> ../common_jars/kafka-log4j-appender-0.10.1.0.jar`

Cloudera Navigator Data Management Release Notes

- kafka-clients-0.10.1.0.jar -> ../common_jars/kafka-clients-0.10.1.0.jar

```
rm kafka-log4j-appender-0.10.1.0.jar
rm kafka-clients-0.10.1.0.jar
```

4. Move the downloaded .jar files to the /usr/share/cmfc/cloudera-navigator-audit-server path.
5. Log in to the Cloudera Manager Admin Console.
6. Restart Navigator Audit Server.

Affects version: Cloudera Manager 5.11.1 and later (Navigator 2.10.1)

Cloudera Issue: NAV-5042

Pending audit events may be lost after role migration

Pending and deleted audit events may be lost during role migration. Pending audit events are those events that have not yet transferred from a source node through the Cloudera Navigator plug-in for the respective role to the Navigator Audit Server. If the Navigator Audit Server is unreachable due to network issues or other issues, the pending audits—which should automatically transfer at the conclusion of the role migration—may be lost.

Workaround: Before migrating a role to a different node:

- Check the state of the Navigator Audit Server (using Cloudera Manager Admin Console).
- Proceed with the role migration only if Navigator Audit Server is running.

Cloudera Issue: NAV-1006

Audit CSV has extra columns and is missing some data

Audit details exported as CSV files do not contain Sentry data. In addition, some of the column names display twice (Operation Text, Database Name, Object Type, for example), although the actual details display only of the duplicate columns.

Workaround: Export audits to JSON format to obtain Sentry data.

Cloudera Issue: NAV-708

Navigator Metadata Server

Navigator Metadata Server fails to start with InvalidClassException, local class incompatible

Navigator Metadata Server fails to start (red status in Cloudera Manager) and the Metadata Server log includes the following exception:

```
2018-09-06 10:38:20,786 WARN jndi [main]:
java.io.InvalidClassException: com.mchange.v2.c3p0.WrapperConnectionPoolDataSource;
local class incompatible: stream classdesc serialVersionUID = 7806429541739165290, local
class serialVersionUID = -7086951306718003710
```

Even if Navigator Metadata Server starts once, it may encounter this problem at a later time, depending on which JAR file is loaded.

Workaround: See Knowledge Base Article [Navigator Metadata Server fails to start with "local class incompatible" error](#)

Affects version: Cloudera Manager version 5.15.1 (Navigator 2.14.1)

Fixed version: Cloudera Manager version 5.15.2 (Navigator 2.14.2)

Cloudera Issue: NAV-6724, TSB-342

Navigator does not mark HDFS entities as deleted when in bulk extraction takes too long to complete

In large HDFS deployments, the `fsimage` takes a long time to index. When an HDFS checkpoint occurs, it creates a new `fsimage`. However if the previous `fsimage` is still in the process of being indexed, Navigator cannot use the incremental changes found in the `inotify` stream because it refers to the newly created `fsimage`.

When this happens, Navigator attempts to start indexing the newer `fsimage`, creating a loop where Navigator can never take advantage of the more efficient change processing through `inotify`. The immediate fallout of this delay is that no HDFS entities deleted in the cluster are marked as deleted in Navigator.

Fixed in version: Cloudera Manager 5.16.1 (2.15.1)

Cloudera Issue: NAV-6456

ClassCastException in Navigator Metadata Server log

When a Hive view is created, then dropped, and then subsequently recreated as a table with the same name as that of the original view, the Hive extraction process shows this exception in Navigator Metadata Server logs:

```
java.lang.ClassCastException: com.cloudera.nav.hive.model.HView cannot be cast to
com.cloudera.nav.hive.model.HTable
```

Affects version: Cloudera Manager 5.14.0 and later (Navigator 2.13.0)

Fixed in version: Cloudera Manager 5.16.1, 5.15.2 (2.15.1, 2.14.2)

Cloudera Issue: NAV-5939

Navigator cannot use PostgreSQL 10.x

Navigator Metadata Server will not work with PostgreSQL versions 10.x. This version incompatibility will result in a message similar to this:

```
2018-08-06 00:00:22,899 WARN com.mchange.v2.async.ThreadPoolAsynchronousRunner [Timer-1]:
com.mchange.v2.async.ThreadPoolAsynchronousRunner$DeadlockDetector@6544a122 - APPARENT
DEADLOCK!!! Creating emergency threads for unassigned pending tasks!
```

Products affected: Cloudera Navigator using PostgreSQL as the database for Navigator Metadata Server

Releases affected: Cloudera Navigator as packaged with Cloudera Manager versions 5.15.0 and 5.15.1.

Users affected: Installations using Cloudera Navigator using PostgreSQL as the database for Navigator Audit Server and Navigator Metadata Server

Severity (Low/Medium/High): High

Impact: Navigator Metadata Server does not function with PostgreSQL database version 10 and above.

Immediate action required: Do not upgrade the Navigator Metadata Server database to PostgreSQL 10.x (allow the Navigator Metadata Server database to remain on a previous working version of PostgreSQL) or upgrade to future Cloudera Manager release where the issue is fixed.

Affects version: Cloudera Manager 5.15.1, 5.15.0 (2.14.1, 2.14.0)

Fixed in version: Cloudera Manager 5.16.1, 5.15.2 (2.15.1, 2.14.2)

Cloudera Issue: DOCS-3686, TSB-334

Navigator Embedded Solr can reach its limit on number of documents it can store

Navigator Metadata Server extracts HDFS entities by performing a one-time bulk extraction and then switching to incremental extraction. In Cloudera Manager releases 5.10.0, 5.10.1 and 5.11.0 (Navigator releases 2.9.0, 2.9.1, and 2.10.0), a problem causes HDFS bulk extraction to be run more than one time, resulting in duplicate relations created

for HDFS. Over time, embedded Solr runs out of document IDs that it can assign to new relations and fails with following error:

```
"Caused by: java.lang.IllegalArgumentException: Too many documents, composite IndexReaders cannot exceed 2147483519"
```

When this happens, Navigator stops any more extraction of data as no new documents can be added to Solr.

The workaround is to upgrade to Cloudera Manager releases 5.10.2, 5.11.1, or 5.12.x (Navigator 2.9.2, 2.10.1, 2.11.x) or later where this issue is fixed. In addition, see [Repairing metadata in the storage directory after upgrading](#).



Important: If you are running Cloudera Navigator in Cloudera Manager 5.10.0, 5.10.1, or 5.11.0, plan your upgrade as soon as possible even if you haven't yet encountered this error. Recovery after encountering the failure increases your risk of losing user-provided metadata. If you aren't able to upgrade immediately, be sure to [backup](#) your Navigator Metadata Server storage directory and database regularly.

Affects version: Cloudera Manager 5.11.0, 5.10.1, 5.10.0 (2.10.0, 2.9.1, 2.9.0)

Fixed in version: Cloudera Manager 5.12.0, 5.11.1, 5.10.2 (Navigator 2.11.0, 2.10.1, 2.9.2)

Cloudera Issue: NAV-5600

"Lineage Diagram is limited to 3000 entities" Error

Navigator caps the number of lineage parents that could be used at one time, however this cap included relations with second level parents (columns). The Navigator Metadata Server log includes an error such as "Lineage Diagram is limited to 3000 entities".

Work-around: Set `nav.capacity.max_nodes_limit=1000` in the Cloudera Manager configuration **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**. Restart Navigator Metadata Server role and check the lineage. The 1000 value assumes that the lineage only touches max of 1000 tables.

Affects version: Cloudera Manager 5.10.0 and later (Navigator 2.9.0)

Fixed in version: Cloudera Manager 5.16.1, 5.15.1, 5.14.4 (Navigator 2.15.1, 2.14.1, 2.13.4)

Cloudera Issue: NAV-6371

Lineage relations are missing for some views with "EndPoint2 must not be null or empty" error

Lineage relations are missing for some views and the Navigator Metadata Server log includes errors with a stack trace that looks something like this:

```
2017-10-25 19:00:40,780 ERROR com.cloudera.nav.persist.impl.CompositeLinker [CDHExecutor-0-CDHUrlClassLoader@381e3]: Internal error while linking. java.lang.IllegalStateException: EndPoint2 must not be null or empty
```

Note that the exception is "EndPoint2 must not be null or empty". If you see errors that don't exactly match this exception, see one of the following similar errors:

- [Log includes the error "EndPoint1 must not be null"](#) on page 2727
- [Log includes the error "EndPoint2 EntityType must not be null unless unlinked"](#) on page 2727

While Navigator is determining lineage relations, it found a relation that was created with the target entity missing, hence Endpoint2 is null. This situation can occur when a Hive view is altered without changing the columns extracted from the source table, such as when the name or other table metadata on the view is changed. For example, the following operation would produce the error:

```
alter view customers_sales_sw set tblproperties("region"="southwest");
```

This problem can occur in Navigator Metadata Server versions 2.9.0, 2.9.1, 2.9.2, 2.10.0, 2.10.1, 2.10.2, 2.11.0, and 2.11.1. These versions correspond to the following Cloudera Manager versions: 5.10.0, 5.10.1, 5.10.2, 5.11.0, 5.11.1, 5.11.2, 5.12.0, and 5.12.1.

Eventually you'll want to upgrade away from an affected Navigator version. If you aren't able to upgrade right away, you can ignore this error with the consequences that your logs are noisy and fill up more quickly. See the Knowledge Base article [EndPoint2 must not be null or empty](#) for more information.

Affects version: Cloudera Manager 5.12.1, 5.11.1 (Navigator 2.11.1, 2.10.1)

Fixed in version: Cloudera Manager 5.13.x, 5.12.2, 5.11.2 (Navigator 2.12.x, 2.11.2, 2.10.2)

Cloudera Issue: NAV-5661, KB-713

Log includes the error "EndPoint1 must not be null"

The following error may appear in the Navigator Metadata Server log:

```
2017-10-17 13:00:23,007 ERROR com.cloudera.nav.hive.extractor.AbstractHiveExtractor
[CDHExecutor-0-CDHUrlClassLoader@14784b7b]: Unable to parse hive view query *: EndPoint1
must not be null or empty
java.lang.IllegalStateException: EndPoint1 must not be null or empty
```

This error occurs because the Hive pull extraction for creating a Hive view produces an incorrect lineage relationship for the Hive view. However, Navigator also receives information for the view creation through the push extractor, which correctly produces the lineage relation. You can safely ignore this error. Note that it is distinct from [Lineage relations are missing for some views with "Endpoint2 must not be null or empty" error](#) on page 2726.

Affects version: Cloudera Manager 5.10.x and later (Navigator 2.9.x)

Cloudera Issue: NAV-4224

Log includes the error "EndPoint2 EntityType must not be null unless unlinked"

The following error may appear in the Navigator Metadata Server log:

```
2017-10-26 15:48:27,440 INFO com.cloudera.nav.hdfs.extractor.HdfsOperationHandler
[CDHExecutor-0-CDHUrlClassLoader@3eca0c78]: Unable to process rename for path
/hbase/data/default/img/f6572262455416bff42f92fd2b0e75c0/.tmp/1cc1af7b0f784572
b821980f6b4c5adc: can't find source information.
2017-10-26 15:48:27,440 ERROR com.cloudera.nav.hdfs.client.InotifyClient
[CDHExecutor-0-CDHUrlClassLoader@3eca0c78]: Error handling event (txid: 2020059072):
Renamed /hbase/data/default/img/f6572262455416bff42f92fd2b0e75c0/.tmp/1cc1af7b0f
784572b821980f6b4c5adc to
/hbase/data/default/img/f6572262455416bff42f92fd2b0e75c0/D/1cc1af7b0f784572b821980f6b4c5adc
at time 1509049550134
java.lang.IllegalStateException: EndPoint2 EntityType must not be null unless unlinked
at com.google.common.base.Preconditions.checkState(Preconditions.java:149)
at com.cloudera.nav.core.model.Relation.validate(Relation.java:287)
```

This error occurs when HDFS files are renamed and the new name is affected by an exclusion filter. The only consequence of the problem is noise in the log. Note that it is distinct from [Lineage relations are missing for some views with "Endpoint2 must not be null or empty" error](#) on page 2726.

This problem prevents new lineage information from being added to existing and new Hive entities.

Affects version: Cloudera Manager 5.7.x and later (Navigator 2.6.x)

Fixed in version: Cloudera Manager 5.14.x (Navigator 2.13.x)

Cloudera Issue: NAV-4654

Table-to-HDFS links not established when Navigator supports multiple, high-availability clusters

When Navigator extracts metadata for multiple clusters and when the clusters are configured for high availability operation, Navigator does not correctly link tables to their HDFS backing data. The result is that lineage between Hive

Cloudera Navigator Data Management Release Notes

tables and their physical data are not created. In addition, some Hive table metadata that is derived from the backing files is not available.

You may see errors in the log such as the following:

```
2018-02-19 09:01:32,999 ERROR com.cloudera.nav.persist.impl.CompositeLinker
[CDHExecutor-0-CDHUrlClassLoader@01010d7e]: Internal error while
linking.java.lang.IllegalArgumentException: expected one element but was:
<com.cloudera.nav.core.model.Source@b22e68ba, com.cloudera.nav.core.model.Source@5b1bac2c>
```

Affects version: Cloudera Manager 5.13.x, 5.12.x (Navigator 2.12.x, 2.11.x)

Fixed in version: Cloudera Manager 5.14.0 (Navigator 2.13.0)

Cloudera Issue: NAV-5749

Purging Deleted Properties shows "Bad Request" error when viewing Hive entity details

After upgrading to Navigator 2.11.0 (Cloudera Manager 5.12.0) and later versions (to Navigator 2.13.0), removing deleted managed metadata properties with the "Purge deleted Properties" command causes a mismatch between the Solr schema and the data stored for entities. (**Administration > Managed Metadata**).

Properties

See the [documentation](#) before creating or modifying properties.

Name	Status	Type	Classes	Multivalued	Actions
Nav_5982	Deleted	Enumeration	Oozie Workflow, Oozie Workflow Instance	No	Actions ▾

The result is that the Navigator console shows the error "Bad Request" when displaying details for Hive entities such as databases, tables, or fields.

When this problem occurs, the Navigator Metadata Server log will include a message such as the following:

```
[CDHExecutor-0-CDHUrlClassLoader@62d310e8]: Internal error while linking.
java.lang.IllegalArgumentException: java.lang.ClassCastException@2f145bbc
```

To avoid this problem, do not purge deleted properties; allow these properties to stay in the list of properties.

Affects version: Cloudera Manager 5.12.0 and later (Navigator 2.11.0)

Fixed in version: Cloudera Manager 5.15.0, 5.14.2, 5.13.3 (Navigator 2.14.0, 2.13.2, 2.12.3)

Cloudera Issue: NAV-5982

Metadata Server log file and spurious error messages

Certain configurations of OS and database (such as PostgreSQL and Ubuntu Linux) may raise spurious error messages regarding non-existent files. Such messages can be safely disregarded, such as this example:

```
Error: [main]: PWC6351: In TLD scanning, the
supplied resource file:/usr/share/java/oracle-connector-java.jar does not
exist.
```

Affects version: Cloudera Manager 5.9.x, 5.8.1 (Navigator 2.8.x, 2.7.1)

Fixed in version: Cloudera Manager 5.10.0 (2.9.0)

Cloudera Issue: NAV-698

Purge

Fail to Purge Hive and Impala Select Operations

In some cases, purge code makes large query requests to solr. In those cases, it can fail with the following error followed by a stack trace:

```
2018-09-22 23:39:55,736 ERROR
com.cloudera.nav.maintenance.purge.hiveandimpala.PurgeHiveOrImpalaSelectOperations
[scheduler_Worker-2]: Failed to purge operations for
DELETE_HIVE_AND_IMPALA_SELECT_OPERATIONS with error
    org.apache.solr.client.solrj.impl.HttpSolrServer$RemoteSolrException: Expected
    mime type application/octet-stream but got application/xml.
```

Work-around: To work around this problem, lower the batch size for Navigator commits to its Solr collection by changing the value of `nav.solr.commit_batch_size` to be lower.

Affects version: Cloudera Manager 5.12.0 (Navigator 2.11) and later

Cloudera Issue: NAV-6789

Purge operation leaves Kite errors in log

Navigator Metadata Server attempts to handle Kite datasets even though Kite datasets are no longer supported. This problem may affect operations such as metadata purge. If you encounter Kite dataset errors such as the following, you need to manually disable Kite dataset extraction.

```
2018-09-23 16:10:43,678 ERROR com.cloudera.nav.hdfs.datasets.KiteDatasetReader
[hdfs-reader-1]: Dataset expected but not found at
/staging/example/.temp/job_1537417909723_0201/mr/job_1537417909723_0201org.kitesdk.data.DatasetNotFoundException:
Descriptor location does not exist
```

To disable Kite dataset extraction, set the following Navigator Metadata Server property in Cloudera Manager, in the **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve) for cloudera-navigator.properties**:

```
nav.extractor.hdfs.datasets.enabled=false
```

Affects version: Cloudera Manager 5.14.0 (Navigator 2.13) and later

Cloudera Issue: NAV-6965

The flag for disabling purge is ignored in the UI

The **Administration > Purge Settings >** page may show scheduled purge operations even if the `nav.purge.enabled` property is set to false. The displayed purge operations will not run.

Affects version: Cloudera Manager 5.14.0, 5.14.1 (Navigator 2.13.1, 2.13.0)

Fixed in version: Cloudera Manager 5.14.2 (Navigator 2.13.2)

Cloudera Issue: NAV-5966

Policy specifications and cluster names affect purge

Policies cannot use cluster names in queries. Cluster name is a derived attribute and cannot be used as-is.

Workaround: When setting move actions for Cloudera Navigator, if there is only one cluster known to the Navigator instance, remove the `clusterName` clause.

If there is more than one cluster known to the Navigator instance, replace `clusterName` with `sourceId`. To get the `sourceId`, issue a query in this format:

```
curl '<nav-url>/api/v9/entities/?query=type%3Asource&limit=100&offset=0'
```

Use the identity of the matching HDFS service for this cluster as the `sourceId`.

Cloudera Issue: NAV-3288

Spark

Spark Lineage Limitations and Requirements

Spark lineage diagrams are supported in the following releases:

- Spark 2.3 as of Cloudera Manager 5.14.0 (Navigator 2.13.0)
- Spark 1.6 as of Cloudera Manager 5.11.0 (Navigator 2.10.0)

Lineage is not available for Spark when Cloudera Manager is running in single user mode. In addition to these requirements, Spark lineage has the following limitations:

- Lineage is produced only for data that is read/written and processed using the Dataframe and SparkSQL APIs. Lineage is not available for data that is read/written or processed using Spark's RDD APIs.
- Lineage information is not produced for calls to aggregation functions such as `groupBy()`.
- The default lineage directory for Spark on Yarn is `/var/log/spark/lineage`. No process or user should write files to this directory—doing so can cause agent failures. In addition, changing the Spark on Yarn lineage directory has no effect: the default remains `/var/log/spark/lineage`.

Cloudera Issue: OPSAPS-39589

Navigator doesn't recognize local files in Spark jobs

Spark jobs can use files on the local filesystem as job inputs or outputs. Navigator, however, only supports HFDS, Hive, and S3 assets as job inputs or outputs. When Navigator extracts metadata from Spark and encounters a local source type, the metadata is discarded and the following error appears in the Navigator Metadata Server log:

```
2018-10-11 12:14:26,192 WARN com.cloudera.nav.api.ApiExceptionMapper
[qtp1574898980-23815]: Unexpected exception.
    java.lang.RuntimeException: Source LOCAL isn't supported for Spark Lineage
```

Affected Versions: Cloudera Manager 5.11 and later (Navigator 2.10)

Fixed in version: Cloudera Manager 5.16.2 (Navigator 2.15.2) Note that the fix does not provide support for extracting metadata for local files for Spark jobs; it prevents the error.

Cloudera Issue: NAV-6811

Too many open files on the Navigator Metadata Server Spark agent

You've encountered this problem if lineage isn't being created for Spark operations and the Cloudera Manager Agent logs include the error:

```
OSError: [Errno 24] Too many open files: '/var/log/spark/lineage'
```

Work-around:

To remove the bottleneck of files and continue processing:

1. From a terminal, login to the host with the issue.
2. [Stop the Cloudera Manager agent.](#)

If the file descriptor limit is hit, the agent is likely to fail to start anyway.

3. Navigate to the Spark lineage directory, by default `/var/log/spark/lineage`.

The directory will contain thousands of files.

4. Move enough files out of this directory to bring count of files within the file descriptor limit.
5. [Start the Cloudera Manager Agent.](#)

The agent will start processing files within few minutes. After successfully processing, the agent will delete the processed files. At the end there will fewer than 10 files present in the directory.

6. Move the files that were moved out back in the directory.

The next set of files will be processed.

Affects version: Cloudera Manager 5.12.1 and later (Navigator 2.11.1)

Fixed in version: Cloudera Manager 5.16.1, 5.15.1, 5.14.4 (Navigator 2.15.1, 2.14.1, 2.13.4)

Cloudera Issue: NAV-6396

Extracting Spark operation produced error message "Solr Manager can be used by only one thread"

When extracting Spark operations, the following error appears in the Navigator Metadata Server log:

```
2018-02-20 09:37:27,869 ERROR com.cloudera.nav.pushextractor.spark.SparkPushExtractor
[qtp937744315-126]: com.cloudera.nav.pushextractor.spark.SparkPushExtractor Error
extracting Spark operation.java.lang.IllegalStateException: Solr Manager can be used by
only one thread.
```

This error does not have consequences for Navigator handling Spark operations.

Affects version: Cloudera Manager 5.14.1, 5.14.0, 5.13.1, 5.13.0, 5.12.x (Navigator 2.13.1, 2.13.0, 2.12.1, 2.12.0, 2.11.x)

Fixed in version: Cloudera Manager 5.15.0, 5.14.2, 5.13.3 (Navigator 2.14.0, 2.13.2, 2.12.3)

Cloudera Issue: NAV-4991

Spark lineage not shown in Navigator for HDFS HA

When HDFS is configured for high availability, Spark lineage doesn't show in Navigator. In these cases, the log includes index out of bounds exceptions such as the following:

```
com.cloudera.nav.pushextractor.spark.SparkPushExtractor Error extracting Spark operation.
java.lang.ArrayIndexOutOfBoundsException
```

Affects version: Cloudera Manager 5.13.0, 5.12.1, 5.12.0, 5.11.x, 5.10.x, 5.9.x (Navigator 2.12.0, 2.11.1, 2.11.0, 2.10.x, 2.9.x, 2.8.x)

Fixed in version: Cloudera Manager 5.14.0, 5.13.1, 5.12.2 (Navigator 2.13.0, 2.12.1, 2.11.2)

Cloudera Issue: NAV-4874

Spark extractor enabled using safety valve deprecated

The Spark extractor included prior to CDH 5.11 and enabled by setting the safety valve, `nav.spark.extraction.enable=true` is being deprecated, and could be removed completely in a future release. If you are upgrading to CDH 5.11 from a deployment that had configured this safety valve, be sure to remove the setting when you upgrade.

Upgrade Issues and Limitations

Issues and limitations matrix by release

Before upgrading Cloudera Navigator, review these version-specific release notes:

Upgrade		Limitations, requirements, preliminary tasks
From...	To...	
2.14 (and lower)	2.15.1 (and higher)	The default filters set for auditing HDFS changed in Cloudera Manager 5.16.1 (Navigator 2.15.1). If you've updated your HDFS audit filters from previous defaults, the new filters are not applied. Whether or not the filters are changed in your upgrade process,

Upgrade		Limitations, requirements, preliminary tasks
From...	To...	
		we recommend that you review your HDFS audit filters to ensure you are collecting the audit events that add value in your environment. See Improvements to Audit Filters for more information.
2.10 (and lower)	2.11.0, 2.11.1, 2.12.0	When using Navigator in multi-cluster environments, avoid upgrading to Cloudera Manager deployments of version 5.12.0, 5.12.1, or 5.13.0 due to a known problem where Navigator does not recognize more than one cluster. Instead, upgrade to Cloudera Manager release 5.12.2, 5.13.1, or 5.14.0 (Navigator 2.11.2, 2.12.1, or 2.13.0). See Cloudera Navigator is not supported in installations in multi-cluster environments for Cloudera Manager versions 5.12.0, 5.12.1, and 5.13.0 on page 2717 for more details. This problem does not affect single cluster deployments or Altus clusters.
2.8 (and lower)	2.9.0, 2.9.1, 2.10.0	Avoid upgrading to Cloudera Manager releases 5.10.0, 5.10.1 and 5.11.0 (Navigator releases 2.9.0, 2.9.1, and 2.10.0) due to the known problem causing the storage directory to fill beyond its capacity. The workaround is to upgrade to Cloudera Manager release 5.10.2, 5.11.1, or 5.12.x (Navigator 2.9.2, 2.10.1, 2.11.x) or later, where this issue is fixed. See Navigator Embedded Solr can reach its limit on number of documents it can store on page 2725 if you are already running one of these releases.
2.10	2.10.1	Upgrading Cloudera Manager 5.11.0 (Cloudera Navigator 2.10) to Cloudera Manager 5.11.1 (Cloudera Navigator 2.10.1) results in failures by Navigator Audit Server to publish to Kafka. See Publishing to Kafka fails after upgrade on page 2723 for details and a workaround.
2.9 (and higher)	2.10 (and higher)	As described later in this table, upgrading across Cloudera Manager 5.10 (Cloudera Navigator 2.9) involved a significant metadata change and therefore the upgrade was very slow. Upgrades between Cloudera Manager 5.10 and later releases do not have a metadata change and the Navigator upgrade is quick.
2.8 (and lower)	2.9 (and higher)	Upgrading to Cloudera Navigator 2.9 and higher (Cloudera Manager 5.10 and higher) can take a significant amount of time, depending on the size of the datadir. Before starting to upgrade to Cloudera Manager 5.10 (which automatically starts the upgrade to Cloudera Navigator 2.9), see Upgrading Cloudera Navigator Can be Extremely Slow on page 2733. Briefly, in this release Solr indexing has been optimized to improve search speed. When the upgrade process completes and Cloudera Navigator services re-start, the Solr indexing upgrade automatically begins. No other actions can be performed until Solr indexing completes (progress message display during this process).
2.6 (and lower)	Any	The Cloudera Navigator Metadata Server requires an upgrade of data in the storage directory. See Upgrading Cloudera Navigator for details.
2.4.0, 2.4.1	2.4.2	After upgrading, requires manual modification to the HDFS extractor state file to change status of UNDELETE taskTypes from SUCCEEDED to FAILED. This manual process is required only for these specific releases listed. See Issues Fixed in Cloudera Navigator 2.4.2 for details.
Any	2.4 – 2.6	Cloudera Navigator 2.4 – 2.6 do not support JDK 1.6, so upgrade any instances of JDK 1.6 to JDK 1.7 or 1.8 before upgrading to these releases. See Step 2: Install Java Development Kit for details.
2.1	2.2	Cloudera Navigator 2.1 used the beta version of the Navigator Metadata Server policy engine. Policies created using that version are not retained during the upgrade.
2.0 (and lower)	2.1 (and higher)	The upgrade wizard for this upgrade path adds the Navigator Metadata Server to the Cloudera Manager cluster. The Navigator Metadata Server is new, and is not the same as the existing Navigator Audit Server database.

Upgrade		Limitations, requirements, preliminary tasks
From...	To...	
1.2	2.0	Cloudera Navigator 1.2 and Cloudera Navigator 2.0 have reached EOL (end of life) status and are no longer supported. There is no wizard for this upgrade path. Cloudera Navigator 2.0 required a clean install. The Navigator 1.2 Navigator Metadata Server was a beta release included with Cloudera Manager 5.0. The 1.2 version of the Navigator Metadata Server role must be removed before the cluster can be upgraded to Cloudera Navigator 2.0 (the roles are not compatible).

Upgrading Cloudera Navigator Can be Extremely Slow

Upgrading a cluster running Cloudera Navigator to Cloudera Manager 5.10 (or higher) can be extremely slow due to an internal change made to the Solr schema in Cloudera Navigator 2.9. A Solr instance is embedded in Cloudera Navigator and supports its Search capabilities. The Solr schema used by Cloudera Navigator has been modified in the 2.9 release to use datatype `long` rather than `string` for an internal `id` field. This change makes Cloudera Navigator far more robust and scalable over the long term.

However, the upgrade process itself can take a significantly long time because the existing Solr documents—the indexed and searchable data structures used by Solr that are contained in the Cloudera Navigator storage directory—are migrated to the new schema. This change to the Solr schema affects only those Cloudera Navigator deployments that use the metadata and lineage features.



Note: Cloudera Navigator deployments that use only the auditing features of the product—not metadata or lineage—are not affected by this issue.

The upgrade process for Cloudera Navigator starts automatically at the end of the Cloudera Manager upgrade, and the migration to the new schema occurs automatically as part of that upgrade process. The amount of time that administrators should allow for this process depends on the quantity stored at the Navigator Metadata Server Storage Dir (`nav.data.dir`, or simply "datadir") location as listed here:

Metadata and lineage usage	Description
None	Deployments that use Cloudera Navigator audit capability only —without metadata or lineage—do not have the issue. Backup the Navigator Metadata Server datadir and then delete it before upgrading.
datadir < 60 GB	Deployments with relatively small Navigator Metadata Server data directories may take 1 to 2 days for the upgrade process to complete. See the Workaround for steps to take before upgrading to Cloudera Manager 5.10 to possibly reduce the upgrade time.
datadir > 60 GB	Deployments with relatively large Navigator Metadata Server data directories may take several days for the upgrade process to complete. See the Workaround for steps to take before upgrading to Cloudera Manager 5.10 to possibly reduce the upgrade time.

Workaround: To reduce the time required for upgrading the Navigator Metadata Server data directories for deployments currently running Cloudera Navigator 2.8 that uses its metadata and lineage features, consider removing unneeded entries from the metadata before the upgrade. The Navigator Purge feature allows you to remove metadata for deleted entities and for entities and operations older than a specified date. For more information on what metadata you can remove with Purge, see [Managing Metadata Storage with Purge](#).

Run purge before starting the Cloudera Manager upgrade (to Cloudera Manager 5.10), following the steps below.



Warning: These steps may mitigate but do not fully resolve the issue. Follow these steps **before** starting the Cloudera Manager upgrade to **Cloudera Manager 5.10** for any Cloudera Manager 5.9 cluster that currently uses the Cloudera Navigator metadata and lineage features.

Cloudera Navigator Data Management Release Notes

- Check the Navigator Metadata Server storage directory size. The path is `/var/lib/cloudera-scm-navigator` (default) unless configured otherwise. If you need to check the setting:
 - Log in to Cloudera Manager Admin Console.
 - Select **Clusters > Cloudera Management Service**.
 - Click **Configuration** and then click the **Navigator Metadata Server** Scope filter:

Navigator Metadata Server	Navigator Metadata Server (node-1)
Storage Dir	<input type="text" value="/var/lib/cloudera-scm-navigator"/>
nav.data.dir	

- Confirm that the cluster uses the default configuration, or make a note of the location specified and the node name.
- Check the size of the actual directory contents. The following example shows a freshly installed system and so it is virtually empty.

```
[root@node-1 ~]# cd /var/lib/cloudera-scm-navigator
[root@node-1 cloudera-scm-navigator]# ls -l
total 12
drwxr-x--- 2 cloudera-scm cloudera-scm 113 Jul 12 06:56 diagnosticData
drwxr-x--- 2 cloudera-scm cloudera-scm 4096 Jul 12 09:16 extractorState
-rw-r----- 1 cloudera-scm cloudera-scm 36 Jul 12 05:54 instance.uuid
drwxr-x--- 4 cloudera-scm cloudera-scm 60 Jul 12 04:18 solr
drwxr-x--- 7 cloudera-scm cloudera-scm 4096 Jul 12 07:26 temp
[root@node-1 cloudera-scm-navigator]# cd solr
[root@node-1 solr]# ls -l
total 4
drwxr-x--- 4 cloudera-scm cloudera-scm 28 Jul 12 05:56 nav_elements
drwxr-x--- 4 cloudera-scm cloudera-scm 28 Jul 12 05:56 nav_relations
-rw-r----- 1 cloudera-scm cloudera-scm 450 Jul 6 16:13 solr.xml
[root@node-1 solr]#
```

- Back up the contents of the directory. Use Cloudera Manager BDR or your preferred method.
- Schedule a purge process as described in [Scheduling the Purge Process](#).



Note: Users and processes cannot access Cloudera Navigator while purge is running.

Set options to purge metadata for deleted HDFS entities and any operations.

- Check the storage directory size again. If needed, re-run the purge with a shorter time span to retain metadata. If the storage directory consumption cannot be reduced below 60GB, do not start the Cloudera Manager upgrade. Instead, contact [Cloudera support](#) to help you with this upgrade.

Cloudera Issue: NAV-5046

Upgrade Limitation (Navigator 1.2 to Navigator 2)

Moving from Cloudera Navigator 1.2 (part of Cloudera Manager 5.0) to Cloudera Navigator 2 requires a fresh install of Cloudera Navigator 2: The beta release of the Navigator Metadata Server included with Cloudera Navigator 1.2 cannot be upgraded. The summary of the process is as follows:

- Delete the Navigator Metadata Server role.
- Remove the content of the Navigator Metadata Server storage directory.
- Add the Navigator Metadata Server role as detailed in [Adding the Navigator Metadata Server Role](#).
- Clear the cache for any browsers that connected to the 1.2 release of Navigator Metadata Server, to avoid errors caused by stale data.

Fixed in version: Navigator 2.0.1

Cloudera Issue: NAV-654

Issues Fixed in Cloudera Navigator Data Management

The following sections describe the issues fixed in each Cloudera Navigator release.

Issues Fixed in Cloudera Navigator 2.15.2

Apache Commons File Upload Vulnerability

The `commons-fileupload` package used by Navigator has been upgraded to 1.3.3 to resolve a security vulnerability indicated by [CVE-2016-3092](#).

Cloudera Issue: NAV-7099

Navigator Metadata Server purge jobs might not run if there were policies configured

Navigator Metadata Server purge could produce messages such as "Checking if maintenance is running" and then failed to run during the available window. This problem occurred when a scheduled purge job waits for extraction tasks to finish but while waiting, a policy job starts, preventing the purge job from running. In this release, policy jobs will not be triggered if purge jobs are scheduled and waiting to start.

Cloudera Issue: NAV-7037

Navigator Metadata Server shows many duplicate lineage relations

Under some circumstances, Navigator Metadata Server created invalid operation to operation execution relations. This caused errors when displaying the lineage in the Navigator console. Also, the metadata purge could not delete the invalid relations. This problem is fixed in this release.

Cloudera Issue: NAV-7035

International characters in tag names

Navigator tags now support UNICODE characters beyond ASCII. Only ASCII text can be used in the name of a user-defined or managed property. Property values can include international characters.

Cloudera Issue: NAV-7011

Reduction in the amount of memory used by Navigator Metadata Server

This release includes changes that improve how Navigator Metadata Server uses memory, reducing the overall heap required.

Cloudera Issue: NAV-6958

Display of purge history was not sorted

The Navigator console content showing the history of purge jobs is now sorted by time in descending order.

Cloudera Issue: NAV-6916

Navigator doesn't recognize local files in Spark jobs

Spark jobs can use files on the local filesystem as job inputs or outputs. Navigator only supports HFDS, Hive, and S3 assets as job inputs or outputs. While local source files are still not supported, this fix produces improves how local source types are handled when extracting metadata from Spark jobs.

Cloudera Issue: NAV-6811

Cloudera Navigator Data Management Release Notes

Agent log files not deleted when Spark1 and Spark2 shared a node

Navigator did not show all lineage information when Spark1 and Spark2 are installed together and Spark roles overlap on a node. Spark agent logs were not processed, potentially leaving thousands of files in the Spark agent logs directory. This problem is fixed in this release.

Cloudera Issue: OPSAPS-49579 and NAV-6662

Navigator Audit Server "CannotSendRequest" error resolved

When an error occurred in the audit pipeline, it was possible that not all components in the Navigator Audit Server recovered fully. One symptom of such a failure is the "CannotSendRequest" error in the Navigator Audit Server log. This problem is fixed in this release.

Cloudera Issue: NAV-6321

Issues Fixed in Cloudera Navigator 2.15.1

Tables that have been deleted in Hive still appear in Navigator

Navigator metadata extraction from Hive uses an extractor state to keep track of the Hive entities committed to the Navigator Metadata Server data storage in Solr. In some cases, such as when the extractor state file is deleted or when Navigator Metadata Server is shut down during extraction from Hive, the extractor state no longer represents the state of the Hive extraction. When this happens, there may be entities listed in Navigator that are not accounted for in the extractor state. When this happens, Navigator no longer updates the Hive entities that are in Solr but not in the extractor state.

This release includes a check to make sure that the extractor state is synchronized with the content of Hive entities in Solr. To enable the check, set the following property in the Cloudera Management Service property **Navigator Metadata Server Advanced Configuration Snippet (Safety Valve)**:

```
nav.extractor.hive.validate_state=true
```

Cloudera Issue: NAV-6745

Navigator Metadata Server purge did not run at the scheduled time

Navigator stored the time for the scheduled purge as a UTC value; however, the value was read by the scheduler as local time. The purge would run later than the expected time by the number of timezones that local time was different from UTC.

Cloudera Issue: NAV-6700

Navigator Metadata Server runs out of memory during purge operations

One of the steps in the process of purging deleted entities from Navigator involves rebuilding lineage relations when one end of the relation included a deleted entity that was to be purged. This process had an issue that produced lineage relations with duplicate endpoint IDs and caused Navigator Metadata Server to stop. The problem has been resolved in this release.

Cloudera Issue: NAV-6675

The Swagger "Try it out" option for the /audits API fails

Using the interactive Navigator API for the `/audits` endpoint failed with the error:

```
No serializer found for class sun.net.www.protocol.http.HttpURLConnection$HttpInputStream and no properties discovered to create BeanSerializer (to avoid exception, disable SerializationConfig.SerializationFeature.FAIL_ON_EMPTY_BEANS) )
```

Also, the documentation accompanying the audit API says a selector is `table` but selector is actually `table_name`.

Workaround: Changing the `limit` to a value less than 10000.

Cloudera Issue: NAV-6664

Support for PostgreSQL 10

Navigator Audit Server and Navigator Metadata Server released in Cloudera Manager 5.15.0 and 5.15.1 did not support PostgreSQL version 10.x.

Cloudera Issue: NAV-6619

Pre-registration was not working for Hive entities

Pre-registration of Hive entities (Htable, HColumn, or HDatabase) fails causing exceptions in the logs such as:

```
java.lang.ClassCastException: com.cloudera.nav.core.model.GenericEntity cannot be cast
to com.cloudera.nav.hive.model.HTable
```

Cloudera Issues: NAV-6608

Navigator Metadata Server has high CPU usage, linking taking too long after Spark extraction

In some cases, Navigator Metadata Server processing requires an extraordinary amount of memory for longer than expected while extracting and linking metadata when Spark extractions are enabled. A symptom of the problem is a flood of Navigator Metadata Server log messages such as:

```
Finished linking relation with id: <id>
```

This change removes duplicate relations from SPARK data flow relations.

Cloudera Issue: NAV-6591

Reduced noise of logging for SparkLinker

The Navigator Metadata Server logs included too much logging detail at the INFO level coming from the linking process that runs against Spark entities. Many of the log entries were moved to DEBUG.

Cloudera Issue: NAV-6560

Upgrading Navigator from Cloudera Manager versions 5.7 and earlier to could lose data when very long text was truncated

Data was truncated if it could not fit field size of 32K characters. This could affect query text among other fields. This fix makes sure that data is preserved.

Cloudera Issue: NAV-6552

Incorrect behavior for lineage from metadata extracted using earlier Navigator releases

When storage directories contain metadata that was extracted before Navigator 2.6, there could be incorrect values for Source Type and Entity Type and missing values for `internalType`, which cause exceptions in the Navigator Metadata Server log and incorrect behavior when creating lineage relations. This change fixes the issue for most cases. Rare cases may require additional metadata changes.

Cloudera Issue: NAV-6514

Navigator does not mark HDFS entities as deleted when in bulk extraction takes too long to complete

In large HDFS deployments, the `fsimage` takes a long time to index. When an HDFS checkpoint occurs, it creates a new `fsimage`. However if the previous `fsimage` is still in the process of being indexed, Navigator cannot use the incremental changes found in the `inotify` stream because it refers to the newly created `fsimage`.

Cloudera Navigator Data Management Release Notes

When this happens, Navigator attempts to start indexing the newer `fsimage`, creating a loop where Navigator can never take advantage of the more efficient change processing through `inotify`. The immediate fallout of this delay is that no HDFS entities deleted in the cluster are marked as deleted in Navigator.

Cloudera Issue: NAV-6456

Missing Audit Events

Audit events were not sorting correctly, causing the audit listing to exclude some events expected in a given time range. This problem is fixed in this release.

Cloudera Issue: NAV-6413

Corrected Metadata Extraction from YARN-run Hive Query

Hive Job Extractor is used to capture Hive queries launched using Hive CLI. The extractor is scheduled by YARN, so the "service type" passed to the extractor is of type YARN. The conflict between the YARN-generated source ID and Hive as the source caused Navigator issues. This change properly manages the Hive CLI query metadata with the YARN service type.

Cloudera Issue: NAV-6409

Too many open files on the Navigator Metadata Server Spark agent

Lineage wasn't being created for Spark operations and the Cloudera Manager Agent logs included the error:

```
OSError: [Errno 24] Too many open files: '/var/log/spark/lineage'
```

Cloudera Issue: NAV-6396

TLS truststore password encrypted

In previous releases, when Cloudera Manager starts Navigator Metadata Server role, it passes trust store password un-encrypted. In this release, when configured to use TLS, the truststore password is encrypted when passed between Cloudera Manager and Navigator.

Cloudera Issue: NAV-6382

"Lineage Diagram is limited to 3000 entities" Error

Navigator capped the number of lineage parents that could be used at one time, however this cap included relations with second level parents (columns). The Navigator Metadata Server log includes an error such as "Lineage Diagram is limited to 3000 entities".

Cloudera Issue: NAV-6371

Malformed requests don't reveal information

Navigator Metadata Server inadvertently disclosed more information than necessary when it responded to malformed requests. In this release, no information is revealed without the call successfully authenticating.

Cloudera Issue: NAV-6346

Navigator Console Handles Increased Property Limits

The Navigator Metadata Server was recently updated to support property names of length up to 500 and no limit on the description field. (See [Name and description lengths increased](#)). In this release, the Navigator console has been modified to handle the increased limits in all contexts.

Cloudera Issue: NAV-6331

Improved Console Editing for Float and Double Property Types

Previously, managed properties defined as `float` or `double` required a decimal point in stored values; this problem is resolved in this release.

Cloudera Issue: NAV-6291

Automated clean-up when Navigator Metadata Server Storage Directory is Cleared

When Navigator Metadata Server role starts, it now checks if the Solr storage directory is empty or not. If it is empty, this check resets the ordinal numbers and starts a new Navigator maintenance history record.

Cloudera Issue: NAV-6267

Linking fails with NullPointerException

Navigator linking failed with errors similar to the following:

```
ERROR org.apache.solr.core.SolrCore [qtp1200906406-61]:
      java.lang.NullPointerException
```

```
ERROR com.cloudera.nav.persist.impl.CompositeLinker
[CDHExecutor-0-CDHUrlClassLoader@23e992bf]: Internal error while linking.
java.lang.NullPointerException
```

Navigator was creating new sources when an exception occurred and it wasn't able to locate the existing source. This release properly handles the original exception and logs the issue clearly.

Cloudera Issue: NAV-6229

Multiple Cluster Extractor Failure

When Navigator Metadata Server is started with a fresh storage directory on Cloudera Manager versions 5.12.0 or later, extractors were not working when a single Cloudera Manager is managing multiple clusters. This problem is fixed in this release.

Cloudera Issue: NAV-6145

Web Server and Platform Version Disclosure Removed

This release includes a change to omit the server and platform versions from the web server response header.

Cloudera Issue: NAV-6019

Security headers improved with Cache-Control

This release includes Cache-Control cookie headers with the setting "no-cache" to ensure that no browser content is cached. Navigator Metadata Server, the web server for the Navigator console, already uses the no-cache Cache-Control header for API interactions. This change extends the security coverage to additional content such as JavaScript.

Cloudera Issue: NAV-6018

ClassCastException in Navigator Metadata Server log

When a Hive view is created, then dropped, and then subsequently recreated as a table with the same name as that of the original view, the Hive extraction process shows this exception in Navigator Metadata Server logs:

```
java.lang.ClassCastException: com.cloudera.nav.hive.model.HView cannot be cast to
com.cloudera.nav.hive.model.HTable
```

Cloudera Issue: NAV-5939

Cloudera Navigator Data Management Release Notes

Less Secure Credentials Protection Policy could expose Azure credentials in audit logs

When you use Cloudera Manager to configure the ADLS Connector service using the Less Secure option for the Credentials Protection Policy, it was possible for Hive audit logs to include Microsoft Azure credentials. If you are using Navigator Audit Server, these credentials could appear in audit reports.

This issue is fixed in this release.

Cloudera Issues: NAV-5861, CDH-56241

API Error when Adding an Audit Filter in Navigator Console (Oracle)

When Navigator Audit Server database is Oracle and when adding Source or Destination filter in the Navigator console Audit page, Navigator returned an API Error. The Navigator Metadata Server log (for the console) included this error:

```
Caused by: java.sql.SQLException: ORA-00932: inconsistent datatypes: expected  
- got CLOB
```

Cloudera Issue: NAV-3582

Issues Fixed in Cloudera Navigator 2.14.2

Connection timeout increased in Navigator Audit Server to avoid broken pipe errors

This release of Navigator increases the connection idle time for the connection between Navigator Audit Server and each Cloudera Manager agent that sends audits to Navigator from audited services. This change will help avoid audit pipeline "broken pipe" errors that can occur when increased load on the audit server node causes Navigator Audit Server not to respond quickly enough.

Cloudera Issue: NAV-6796

Navigator Metadata Server failed to start with `InvalidClassException`, local class incompatible

Navigator Metadata Server failed to start (red status in Cloudera Manager) and the Metadata Server log included the following exception:

```
2018-09-06 10:38:20,786 WARN jndi [main]:  
java.io.InvalidClassException: com.mchange.v2.c3p0.WrapperConnectionPoolDataSource;  
local class incompatible: stream classdesc serialVersionUID = 7806429541739165290, local  
class serialVersionUID = -7086951306718003710
```

Even if Navigator Metadata Server started once, it may encounter this problem at a later time, depending on which JAR file is loaded.

Workaround: See Knowledge Base Article [Navigator Metadata Server fails to start with "local class incompatible" error](#)

Affects version: Cloudera Manager version 5.15.1 (Navigator 2.14.1)

Cloudera Issue: NAV-6724

Navigator Metadata Server purge did not run at the scheduled time

Navigator stored the time for the scheduled purge as a UTC value; however, the value was read by the scheduler as local time. The purge would run later than the expected time by the number of timezones that local time was different from UTC.

Cloudera Issue: NAV-6700

Support for PostgreSQL 10

Navigator Audit Server and Navigator Metadata Server released in Cloudera Manager 5.15.0 and 5.15.1 did not support PostgreSQL version 10.x.

Cloudera Issue: NAV-6619

Navigator Metadata Server has high CPU usage, linking taking too long after Spark extraction

In some cases, Navigator Metadata Server processing requires an extraordinary amount of memory for longer than expected while extracting and linking metadata when Spark (Spark 1 or Spark 2) extractions are enabled. A symptom of the problem is a flood of Navigator Metadata Server log messages such as:

```
Finished linking relation with id: <id>
```

This change removes duplicate relations from SPARK data flow relations.

Cloudera Issue: NAV-6591

ClassCastException in Navigator Metadata Server log

When a Hive view is created, then dropped, and then subsequently recreated as a table with the same name as that of the original view, the Hive extraction process shows this exception in Navigator Metadata Server logs:

```
java.lang.ClassCastException: com.cloudera.nav.hive.model.HView cannot be cast to
com.cloudera.nav.hive.model.HTable
```

Cloudera Issue: NAV-5939

Authentication fails with javax.naming.PartialResultException when using LDAP

Previously, if Navigator Metadata Server is configured with **External Authentication Type** "LDAP" and the LDAP server is Active Directory, when the **LDAP Group Search Base** property is configured with the root of the directory tree, authentication fails in the UI with `Incorrect username or password`.

Navigator Metadata Server logs show the following exception:

```
2018-10-24 12:21:04,619 ERROR com.cloudera.nav.auth.DelegatingNavAuthProvider
[qtpl989132530-59]: External Authentication Service threw exception during authentication
process.
org.springframework.ldap.PartialResultException: Unprocessed Continuation Reference(s);
nested exception is javax.naming.PartialResultException: Unprocessed Continuation
Reference(s); remaining name 'DC=ad,DC=sec,DC=example,DC=com'
```

There were two common workarounds for this problem: use the Global Catalog or configure a more specific location in the directory tree for the **LDAP Group Search Base**. In this release, the root-level **LDAP Group Search Base** is correctly evaluated and no work-around is required. If you have configured one of these workarounds, you can keep the configuration as is with no change of behavior.

Affected version: Cloudera Manager 5.13.x (Navigator 2.12.x)

Cloudera Issue: NAV-5669

Issues Fixed in Cloudera Navigator 2.14.1

"GenericEntity cannot be cast" Exception

When storage directories contain metadata that was extracted before Navigator 2.6, there could be incorrect values for `Source Type` and `Entity Type` and missing values for `internalType`, which cause exceptions in the Navigator Metadata Server log and incorrect behavior when creating lineage relations. Exceptions were similar to the following:

```
java.lang.ClassCastException: com.cloudera.nav.core.model.GenericEntity cannot be cast to com.cloudera.nav.hive.model.HQuery
```

```
java.lang.ClassCastException: com.cloudera.nav.core.model.GenericEntity cannot be cast to com.cloudera.nav.impala.model.ImpalaQuery
```

This change fixes the issue for most of the cases. Rare cases may require additional metadata changes.

Cloudera Issue: NAV-6514

Missing Audit Events

Audit events were not sorting correctly, causing the audit listing to exclude some events expected in a given time range. This problem is fixed in this release.

Cloudera Issue: NAV-6413

Corrected Metadata Extraction from YARN-run Hive Query

Hive Job Extractor is used to capture Hive queries launched using Hive CLI. The extractor is scheduled by YARN, so the "service type" passed to the extractor is of type YARN. The conflict between the YARN-generated source ID and Hive as the source caused Navigator issues. This change properly manages the Hive CLI query metadata with the YARN service type.

Cloudera Issue: NAV-6409

Too many open files on the Navigator Metadata Server Spark agent

Lineage wasn't being created for Spark operations and the Cloudera Manager Agent logs included the error:

```
OSError: [Errno 24] Too many open files: '/var/log/spark/lineage'
```

Cloudera Issue: NAV-6396

Trust Store password encrypted

In previous releases, when Cloudera Manager starts Navigator Metadata Server role, it passed trust store passwords un-encrypted. In this release, when configured to use TLS, the trust store password is encrypted when passed between Cloudera Manager and Navigator.

"Lineage Diagram is limited to 3000 entities" Error

Navigator capped the number of lineage parents that could be used at one time, however this cap included relations with second level parents (columns). The Navigator Metadata Server log includes an error such as "Lineage Diagram is limited to 3000 entities".

Cloudera Issue: NAV-6371

Navigator Console Handles Increased Property Limits

The Navigator Metadata Server was recently updated to support property names of length up to 500 and no limit on the description field. (See [Name and description lengths increased](#)). In this release, the Navigator console has been modified to handle the increased limits in all contexts.

Cloudera Issue: NAV-6331

Improved Console Editing for Float and Double Property Types

Previously, managed properties defined as `float` or `double` required a decimal point in stored values; this problem is resolved in this release.

Cloudera Issue: NAV-6291

Automated clean-up when Navigator Metadata Server Storage Directory is Cleared

When Navigator Metadata Server role starts, it now checks if the Solr storage directory is empty or not. If it is empty, this check resets the ordinal numbers and starts a new Navigator maintenance history record.

Cloudera Issue: NAV-6267

Linking fails with NullPointerException

Navigator linking failed with errors similar to the following:

```
ERROR org.apache.solr.core.SolrCore [qtp1200906406-61]:
    java.lang.NullPointerException
```

```
ERROR com.cloudera.nav.persist.impl.CompositeLinker
[CDHExecutor-0-CDHUrlClassLoader@23e992bf]: Internal error while linking.
java.lang.NullPointerException
```

Navigator was creating new sources when an exception occurred and it wasn't able to locate the existing source. This release properly handles the original exception and logs the issue clearly.

Cloudera Issue: NAV-6229

Multiple Cluster Extractor Failure

When Navigator Metadata Server is started with a fresh storage directory on Cloudera Manager versions 5.12.0 or later, extractors were not working when a single Cloudera Manager is managing multiple clusters. This problem is fixed in this release.

Cloudera Issue: NAV-6145

HDFS metadata extraction performance improvement

Navigator processing for HDFS metadata extraction can take long enough that Navigator cannot finish indexing the `fsimage` before the edit log is checkpointed into a new `fsimage`. This causes extraction to get stuck in a loop of always parsing the `fsimage` in bulk extraction mode (which is slow) rather than moving into incremental extraction from the edit log. This release includes changes that allow Navigator to avoid the bulk extraction loop, therefore reducing the duration of HDFS extraction by as much as half for the initial extraction; performance of subsequent extractions can improve significantly, benefiting both by moving to incremental extraction and by improvements in this fix.

Cloudera Issue: NAV-6033

Web Server and Platform Version Disclosure Removed

This release includes a change to omit the server and platform versions from the web server response header.

Cloudera Issue: NAV-6019

Cloudera Navigator Data Management Release Notes

Pre-registration of Hive/Impala Queries

Previously, it was not possible to pre-register data asset entities by calling the POST /entities API. This change allows application developers to pre-register operations by exposing the method used to generate the operation source IDs.

The Navigator SDK includes an example of pre-registering entities in the `CustomHiveLineageCreator.java` source file.

Cloudera Issue: NAV-5940

Less Secure Credentials Protection Policy could expose Azure credentials in audit logs

When you use Cloudera Manager to configure the ADLS Connector service using the Less Secure option for the Credentials Protection Policy, it was possible for Hive audit logs to include Microsoft Azure credentials. If you are using Navigator Audit Server, these credentials could appear in audit reports.

This issue is fixed in this release.

Cloudera Issues: NAV-5861, CDH-56241

API Error when Adding an Audit Filter in Navigator Console (Oracle)

When Navigator Audit Server database is Oracle and when adding filters that use Source, Destination, Operation Text, Resource Path, or Status properties in the Navigator console Audit page, Navigator returned an API Error. The Navigator Metadata Server log (for the console) included this error:

```
Caused by: java.sql.SQLException: ORA-00932: inconsistent datatypes: expected  
- got CLOB
```

This problem is fixed in this release.

Cloudera Issue: NAV-3582

Issues Fixed in Cloudera Navigator 2.14.0

Console doesn't come up, NAV_EXTRACTOR_STATUS doesn't exist

Libraries included in the Navigator release for Cloudera Manager 5.13 were named such that the upgrade marked the schema with the wrong Navigator version. The result is that schema changes between the original version and the final upgrade version are not run. If your Navigator upgrade encounters this problem, the Navigator console does not come up and you'll see the following error in the Navigator Metadata Server log:

```
Table 'nav_meta_user.NAV_EXTRACTOR_STATUS' doesn't exist
```

Cloudera Issue: NAV-6213

Error dialog block developer mode on IE 11

When using Microsoft Internet Explorer version 11 with developer-mode turned on, an error dialog appeared in the Navigator console that made it difficult to use the console. The error also appeared in the IE 11 developer console.

Cloudera Issue: NAV-6181

Extractors not working in multi-cluster environment

When Navigator Metadata Server is started with a fresh storage directory on versions 5.12.0 or later, extractors were not working when a single Cloudera Manager is managing multiple clusters.

Cloudera Issue: NAV-6145

Security headers improved with Cache-Control

This release includes Cache-Control cookie headers with the setting "no-cache" to ensure that no browser content is cached. Navigator Metadata Server, the web server for the Navigator console, already uses the no-cache Cache-Control header for API interactions. This change extends the security coverage to additional content such as JavaScript.

Cloudera Issue: NAV-6018

Purging Deleted Properties shows "Bad Request" error when viewing Hive entity details

After upgrading to Navigator 2.11.0 (Cloudera Manager 5.12.0) and later versions (to Navigator 2.13.0), removing deleted managed metadata properties with the "Purge deleted Properties" command caused a mismatch between the Solr schema and the data stored for entities. (**Administration > Managed Metadata**).

The result was that the Navigator console shows the error "Bad Request" when displaying details for Hive entities such as databases, tables, or fields.

When this problem occurs, the Navigator Metadata Server log included a message such as the following:

```
[CDHExecutor-0-CDHUrlClassLoader@62d310e8]: Internal error while linking.  
java.lang.IllegalArgumentException: java.lang.ClassCastException@2f145bbc
```

Cloudera Issue: NAV-5982

Console description improved for Purge cutoff time

The Navigator Purge Settings page indicated that the purge operation would end at the time specified. However, the specified time is used to stop *new* purge tasks from starting. *Currently running* purge tasks may continue to run after the specified end time. The text in the console has been changed to properly reflect the fact that purge tasks may continue to run past the time limit.

Cloudera Issue: NAV-5967

Managed metadata property values hidden

Property values were not visible when the property is defined with multivalued enabled and the type was any numeric type.

Cloudera Issue: NAV-5899

Less Secure Credentials Protection Policy could expose Azure credentials in audit logs

When you use Cloudera Manager to configure the ADLS Connector service using the Less Secure option for the Credentials Protection Policy, it was possible for Hive audit logs to include Microsoft Azure credentials. If you are using Navigator Audit Server, these credentials could appear in audit reports.

This issue is fixed in this release.

Cloudera Issues: NAV-5861, CDH-56241

Multi-cluster support for Hive lineage

When there were multiple Hive MetaStore (HMS) sources associated with a single Hive service, the method Navigator used to determine the source was not accurate, resulting in Impala and Spark extractors creating lineage relations that referred to the wrong sources and cross-cluster lineage relations were not generated.

Cloudera Issue: NAV-5859

Cloudera Navigator Data Management Release Notes

Extracting Spark operation produced error message "Solr Manager can be used by only one thread"

When extracting Spark operations, the following error appeared in the Navigator Metadata Server log:

```
2018-02-20 09:37:27,869 ERROR com.cloudera.nav.pushextractor.spark.SparkPushExtractor [qtp937744315-126]: com.cloudera.nav.pushextractor.spark.SparkPushExtractor Error extracting Spark operation.java.lang.IllegalStateException: Solr Manager can be used by only one thread.
```

This error does not have consequences for Navigator handling Spark operations.

Cloudera Issue: NAV-4991

Issues Fixed in Cloudera Navigator 2.13.4

"GenericEntity cannot be cast" Exception

When storage directories contain metadata that was extracted before Navigator 2.6, there could be incorrect values for `Source Type` and `Entity Type` and missing values for `internalType`, which cause exceptions in the Navigator Metadata Server log and incorrect behavior when creating lineage relations. Exceptions were similar to the following:

```
java.lang.ClassCastException: com.cloudera.nav.core.model.GenericEntity cannot be cast to com.cloudera.nav.hive.model.HQuery
```

```
java.lang.ClassCastException: com.cloudera.nav.core.model.GenericEntity cannot be cast to com.cloudera.nav.impala.model.ImpalaQuery
```

This change fixes the issue for most of the cases. Rare cases may require additional metadata changes.

Cloudera Issue: NAV-6514

Missing Audit Events

Audit events were not sorting correctly, causing the audit listing to exclude some events expected in a given time range. This problem is fixed in this release.

Cloudera Issue: NAV-6413

Corrected Metadata Extraction from YARN-run Hive Query

Hive Job Extractor is used to capture Hive queries launched using Hive CLI. The extractor is scheduled by YARN, so the "service type" passed to the extractor is of type YARN. The conflict between the YARN-generated source ID and Hive as the source caused Navigator issues. This change properly manages the Hive CLI query metadata with the YARN service type.

Cloudera Issue: NAV-6409

Too many open files on the Navigator Metadata Server Spark agent

Lineage wasn't being created for Spark operations and the Cloudera Manager Agent logs included the error:

```
OSError: [Errno 24] Too many open files: '/var/log/spark/lineage'
```

Cloudera Issue: NAV-6396

"Lineage Diagram is limited to 3000 entities" Error

Navigator capped the number of lineage parents that could be used at one time, however this cap included relations with second level parents (columns). The Navigator Metadata Server log includes an error such as "Lineage Diagram is limited to 3000 entities".

Cloudera Issue: NAV-6371**Navigator Console Handles Increased Property Limits**

The Navigator Metadata Server was recently updated to support property names of length up to 500 and no limit on the description field. (See [Name and description lengths increased](#)). In this release, the Navigator console has been modified to handle the increased limits in all contexts.

Cloudera Issue: NAV-6331**Improved Console Editing for Float and Double Property Types**

Previously, managed properties defined as `float` or `double` required a decimal point in stored values; this problem is resolved in this release.

Cloudera Issue: NAV-6291**Automated clean-up when Navigator Metadata Server Storage Directory is Cleared**

When Navigator Metadata Server role starts, it now checks if the Solr storage directory is empty or not. If it is empty, this check resets the ordinal numbers and starts a new Navigator maintenance history record.

Cloudera Issue: NAV-6267**Linking fails with NullPointerException**

Navigator linking failed with errors similar to the following:

```
ERROR org.apache.solr.core.SolrCore [qtp1200906406-61]:
    java.lang.NullPointerException
```

```
ERROR com.cloudera.nav.persist.impl.CompositeLinker
[CDHExecutor-0-CDHUrlClassLoader@23e992bf]: Internal error while linking.
java.lang.NullPointerException
```

Navigator was creating new sources when an exception occurred and it wasn't able to locate the existing source. This release properly handles the original exception and logs the issue clearly.

Cloudera Issue: NAV-6229**Multiple Cluster Extractor Failure**

When Navigator Metadata Server is started with a fresh storage directory on Cloudera Manager versions 5.12.0 or later, extractors were not working when a single Cloudera Manager is managing multiple clusters. This problem is fixed in this release.

Cloudera Issue: NAV-6145**Web Server and Platform Version Disclosure Removed**

This release includes a change to omit the server and platform versions from the web server response header.

Cloudera Issue: NAV-6019**Pre-registration of Hive/Impala Queries**

Previously, it was not possible to pre-register data asset entities by calling the `POST /entities` API. This change allows application developers to pre-register operations by exposing the method used to generate the operation source IDs.

The Navigator SDK includes an example of pre-registering entities in the `CustomHiveLineageCreator.java` source file.

Cloudera Issue: NAV-5940

Issues Fixed in Cloudera Navigator 2.13.3

There were no issues fixed in Cloudera Navigator 2.13.3.

Issues Fixed in Cloudera Navigator 2.13.2

Purging Deleted Properties shows "Bad Request" error when viewing Hive entity details

After upgrading to Navigator 2.11.0 (Cloudera Manager 5.12.0) and later versions (to Navigator 2.13.0), removing deleted managed metadata properties with the "Purge deleted Properties" command caused a mismatch between the Solr schema and the data stored for entities. (**Administration > Managed Metadata**).

The result was that the Navigator console shows the error "Bad Request" when displaying details for Hive entities such as databases, tables, or fields.

When this problem occurs, the Navigator Metadata Server log included a message such as the following:

```
[CDHExecutor-0-CDHUrlClassLoader@62d310e8]: Internal error while linking.  
java.lang.IllegalArgumentException: java.lang.ClassCastException@2f145bbc
```

Cloudera Issue: NAV-5982

Console description improved for Purge cutoff time

The Navigator Purge Settings page indicated that the purge operation would end at the time specified. However, the specified time is used to stop *new* purge tasks from starting. *Currently running* purge tasks may continue to run after the specified end time. The text in the console has been changed to properly reflect the fact that purge tasks may continue to run past the time limit.

Cloudera Issue: NAV-5967

The flag for disabling purge was ignored in the UI

The **Administration > Purge Settings** page could show scheduled purge operations even if the `nav.purge.enabled` property was set to false. The displayed purge operations would not run.

This issue has been resolved.

Cloudera Issue: NAV-5966

Managed metadata property values hidden

Property values were not visible when the property is defined with multivalued enabled and the type was any numeric type.

Cloudera Issue: NAV-5899

Less Secure Credentials Protection Policy could expose Azure credentials in audit logs

When you use Cloudera Manager to configure the ADLS Connector service using the Less Secure option for the Credentials Protection Policy, it was possible for Hive audit logs to include Microsoft Azure credentials. If you are using Navigator Audit Server, these credentials could appear in audit reports.

This issue is fixed in this release.

Cloudera Issues: NAV-5861, CDH-56241

Missing Impala audit events for database-level operations

Impala audit events at the database level were incorrectly recorded, making it seem like they are missing from the Navigator audit data.

This issue is fixed in this release.

Cloudera Issue: NAV-5648

Extracting Spark operation produced error message "Solr Manager can be used by only one thread"

When extracting Spark operations, the following error appeared in the Navigator Metadata Server log:

```
2018-02-20 09:37:27,869 ERROR com.cloudera.nav.pushextractor.spark.SparkPushExtractor
[qtp937744315-126]: com.cloudera.nav.pushextractor.spark.SparkPushExtractor Error
extracting Spark operation.java.lang.IllegalStateException: Solr Manager can be used by
only one thread.
```

This error does not have consequences for Navigator handling Spark operations.

Cloudera Issue: NAV-4991**Issues Fixed in Cloudera Navigator 2.13.1**

There were no issues fixed in Cloudera Navigator 2.13.1.

Issues Fixed in Cloudera Navigator 2.13.0

Log includes the error "EndPoint2 EntityType must not be null unless unlinked"

The following error appeared in the Navigator Metadata Server log:

```
2017-10-26 15:48:27,440 INFO com.cloudera.nav.hdfs.extractor.HdfsOperationHandler
[CDHExecutor-0-CDHUrlClassLoader@3eca0c78]: Unable to process rename for path
/hbase/data/default/img/f6572262455416bff42f92fd2b0e75c0/.tmp/1cc1af7b0f784572
b821980f6b4c5adc: can't find source information.
2017-10-26 15:48:27,440 ERROR com.cloudera.nav.hdfs.client.InotifyClient
[CDHExecutor-0-CDHUrlClassLoader@3eca0c78]: Error handling event (txid: 2020059072):
Renamed /hbase/data/default/img/f6572262455416bff42f92fd2b0e75c0/.tmp/1cc1af7b0f
784572b821980f6b4c5adc to
/hbase/data/default/img/f6572262455416bff42f92fd2b0e75c0/D/1cc1af7b0f784572b821980f6b4c5adc
at time 1509049550134
java.lang.IllegalStateException: EndPoint2 EntityType must not be null unless unlinked
at com.google.common.base.Preconditions.checkState(Preconditions.java:149)
at com.cloudera.nav.core.model.Relation.validate(Relation.java:287)
```

This error occurred when HDFS files are renamed and the new name is affected by an exclusion filter. The only consequence of the problem is noise in the log.

If you are not able to upgrade, you can ignore the error.

Cloudera Issue: NAV-4654**Multi-cluster support**

Navigator was not able to uniquely identify more than one cluster in installations of Cloudera Manager that supported multiple clusters. This problem only occurs for Cloudera Manager deployments of version 5.12.0, 5.12.1 or 5.13.0. This problem does not affect single cluster deployments or Altus clusters.

This issue has been resolved. The upgrade process to this release of Cloudera Manager resolves the problem.



Note: If your multi-cluster installation is configured for high-availability, you will see exceptions thrown by Navigator while linking HDFS entities. This is a known issue for Cloudera Manager 5.12.0, 5.12.1, 5.13.0 and 5.13.1 releases.

Cloudera Issue: NAV-5615

Cloudera Navigator Data Management Release Notes

Table-to-HDFS links not established when Navigator supported multiple, high-availability clusters

When Navigator extracted metadata for multiple clusters and when the clusters were configured for high availability operation, Navigator did not correctly link tables to their HDFS backing data. The result was that lineage between Hive tables and their physical data were not created. In addition, some Hive table metadata that was derived from the backing files was not available.

You could see errors in the log such as the following:

```
2018-02-19 09:01:32,999 ERROR com.cloudera.nav.persist.impl.CompositeLinker
[CDHExecutor-0-CDHUrlClassLoader@01010d7e]: Internal error while
linking.java.lang.IllegalArgumentException: expected one element but was:
<com.cloudera.nav.core.model.Source@b22e68ba, com.cloudera.nav.core.model.Source@5b1bac2c>
```

This issue has been resolved.

Cloudera Issue: NAV-5749

Difficult to know how long Navigator upgrade would take

The upgrade process for Navigator now includes messages to show the estimated time remaining at each upgrade step.

Cloudera Issue: NAV-5639

Not all services were included in log messages during purge operation

The log messages produced during a purge operation are now consistent across sources; they all include counts for the number of entities removed.

Cloudera Issue: NAV-3389

Technical metadata for Hive table partitions was cut off

This release provides a "Show more" link in table detail pages when the list of partition columns doesn't fit in the window.

Cloudera Issue: NAV-4840

Audits Dashboard showed incorrect time period

In certain cases the dashboard link for "Accesses Denied" would not open the correct time period on the Audits page.

This issue has been resolved.

Cloudera Issue: NAV-5689

Duplicate values could be added to dynamic facets

Values added to a dynamic facet were not properly validated. This release includes validation so duplicate values cannot be included in a facet.

Cloudera Issue: NAV-5572

Spark lineage not shown in Navigator for HDFS HA

When HDFS is configured for high availability, Spark lineage didn't show in Navigator. In these cases, the log includes index out of bounds exceptions such as the following:

```
com.cloudera.nav.pushextractor.spark.SparkPushExtractor Error extracting Spark operation.
java.lang.ArrayIndexOutOfBoundsException
```

This issue has been resolved.

Cloudera Issue: NAV-4874

Issues Fixed in Cloudera Navigator 2.12.3

Purging Deleted Properties shows "Bad Request" error when viewing Hive entity details

After upgrading to Navigator 2.11.0 (Cloudera Manager 5.12.0) and later versions (to Navigator 2.13.0), removing deleted managed metadata properties with the "Purge deleted Properties" command caused a mismatch between the Solr schema and the data stored for entities. (**Administration > Managed Metadata**).

The result was that the Navigator console shows the error "Bad Request" when displaying details for Hive entities such as databases, tables, or fields.

When this problem occurs, the Navigator Metadata Server log included a message such as the following:

```
[CDHExecutor-0-CDHUrlClassLoader@62d310e8]: Internal error while linking.
java.lang.IllegalArgumentException: java.lang.ClassCastException@2f145bbc
```

Cloudera Issue: NAV-5982

Extracting Spark operation produced error message "Solr Manager can be used by only one thread"

When extracting Spark operations, the following error appeared in the Navigator Metadata Server log:

```
2018-02-20 09:37:27,869 ERROR com.cloudera.nav.pushextractor.spark.SparkPushExtractor
[qtp937744315-126]: com.cloudera.nav.pushextractor.spark.SparkPushExtractor Error
extracting Spark operation.java.lang.IllegalStateException: Solr Manager can be used by
only one thread.
```

This error does not have consequences for Navigator handling Spark operations.

Cloudera Issue: NAV-4991

Issues Fixed in Cloudera Navigator 2.12.2

Missing Impala audit events for database-level operations

Impala audit events at the database level were incorrectly recorded, making it seem like they were missing from the Navigator audit data.

This issue has been resolved.

Cloudera Issue: NAV-5648

Issues Fixed in Cloudera Navigator 2.12.1

Tags disappeared on secure cluster when a change caused HDFS bulk extraction to run a second time.

When Navigator HDFS extraction occurs for the first time, Navigator performs a bulk extraction. Subsequent bulk extractions should refresh only the changed metadata. However, in previous releases, Navigator extraction refreshed all metadata rather than building on the existing metadata. The result was that tags added before the first bulk extraction were missing after the second bulk extraction. This problem only occurs when some change happens on the cluster to trigger the second bulk extraction, such as resetting the HDFS `inotify` stream in an Isilon-backed cluster.

This issue has been resolved.

Cloudera Issue: NAV-5620

Cloudera Navigator Data Management Release Notes

Multi-cluster support

Navigator was not able to uniquely identify more than one cluster in installations of Cloudera Manager that supported multiple clusters. This problem only occurs for Cloudera Manager deployments of version 5.12.0, 5.12.1 or 5.13.0. This problem does not affect single cluster deployments or Altus clusters.

This issue has been resolved. The upgrade process to this release of Cloudera Manager resolves the problem.



Note: If your multi-cluster installation is configured for high-availability, you will see exceptions thrown by Navigator while linking HDFS entities. This is a known issue for Cloudera Manager 5.12.0, 5.12.1, 5.13.0 and 5.13.1 releases.

Cloudera Issue: NAV-5615

Navigator did not show metadata content for Isilon

Navigator failed to show file system results in a secure cluster backed by a non-HDFS file system, such as EMC Isilon.

This issue has been resolved.

Cloudera Issue: NAV-5375

Incorrect operation of the Search filter for "Deleted Time"

The Deleted Time search filter did not produce expected results for filtering entities.

This issue has been resolved.

Cloudera Issue: NAV-1168

Support for multiple metadata filters with the same name, different namespaces

Multiple managed metadata filters with same name but different namespaces could not be added as filters to search.

This issue has been resolved.

Cloudera Issue: NAV-5429

User-defined relations did not highlight appropriately in lineage diagrams

In Navigator lineage diagrams, when you highlight a path, all relations that aren't in the highlighted path fade. However, in previous releases, user-defined relations in lineage pictures did not fade when they were not in the highlighted path.

This issue has been resolved.

Cloudera Issue: NAV-5574

Spark lineage not shown in Navigator for HDFS HA

When HDFS is configured for high availability, Spark lineage didn't show in Navigator. In these cases, the log includes index out of bounds exceptions such as the following:

```
com.cloudera.nav.pushextractor.spark.SparkPushExtractor Error extracting Spark operation.  
java.lang.ArrayIndexOutOfBoundsException
```

This issue has been resolved.

Cloudera Issue: NAV-4874

Issues Fixed in Cloudera Navigator 2.12.0

Lineage relations were missing for some views with "Endpoint2 must not be null or empty" error

Lineage relations were missing for some views and the Navigator Metadata Server log includes errors with a stack trace that looks similar to the following:

```
2017-10-25 19:00:40,780 ERROR com.cloudera.nav.persist.impl.CompositeLinker
[CDHExecutor-0-CDHUrlClassLoader@381e3]: Internal error while linking.
java.lang.IllegalStateException: EndPoint2 must not be null or empty
```

While Navigator determined lineage relations, it found a relation that was created with the target entity missing, hence "Endpoint2 is null." This situation can occur when a Hive view is altered without changing the columns extracted from the source table, such as when the name or other table metadata on the view is changed. For example, the following operation would produce the error:

```
alter view customers_sales_sw set tblproperties("region"="southwest");
```

This problem can occur in Navigator Metadata Server versions 2.9.0, 2.9.1, 2.9.2, 2.10.0, 2.10.1, 2.10.2, 2.11.0, and 2.11.1. These versions correspond to the following Cloudera Manager versions: 5.10.0, 5.10.1, 5.10.2, 5.11.0, 5.11.1, 5.11.2, 5.12.0, and 5.12.1.

This problem prevents new lineage information from being added to existing and new Hive entities. See the Knowledge Base article [EndPoint2 must not be null or empty](#) for more information.

Cloudera Issue: NAV-5661

Limitation on editing certain managed metadata properties

Managed metadata properties (without class associations) that were added using the Navigator API could not be edited after the fact through the Cloudera Navigator console.

This issue has been resolved.

Cloudera Issue: NAV-5363

Hue links were missing from Navigator console

The links to Hue from entities in the Navigator console were not generated. Each time the console attempted to show a link, the Navigator Metadata Server log included an error similar to the following:

```
10:05:32.539 AM ERROR HueUtility
[qtp778450792-57]: Unable to get Hue link from CM.
java.lang.LinkageError: loader constraint violation: when resolving method
"com.cloudera.nav.server.NavOptions.getPreferredBaseHueURLForCluster(Ljava/lang/String;)Lcom/google/common/base/Optional;"
the class loader (instance of org/eclipse/jetty/webapp/WebAppClassLoader) of the current
class, com/cloudera/nav/utils/HueUtility, and the class loader (instance of
sun/misc/Launcher$AppClassLoader) for resolved class, com/cloudera/nav/server/NavOptions,
have different Class objects for the type
ter(Ljava/lang/String;)Lcom/google/common/base/Optional; used in the signature{
```

This issue has been resolved.

Cloudera Issue: NAV-4756

Issues Fixed in Cloudera Navigator 2.11.2

Lineage relations were missing for some views with "Endpoint2 must not be null or empty" error

Lineage relations were missing for some views and the Navigator Metadata Server log includes errors with a stack trace that looks similar to the following:

```
2017-10-25 19:00:40,780 ERROR com.cloudera.nav.persist.impl.CompositeLinker
[CDHExecutor-0-CDHUrlClassLoader@381e3]: Internal error while linking.
java.lang.IllegalStateException: EndPoint2 must not be null or empty
```

While Navigator determined lineage relations, it found a relation that was created with the target entity missing, hence "Endpoint2 is null." This situation can occur when a Hive view is altered without changing the columns extracted from the source table, such as when the name or other table metadata on the view is changed. For example, the following operation would produce the error:

```
alter view customers_sales_sw set tblproperties("region"="southwest");
```

This problem can occur in Navigator Metadata Server versions 2.9.0, 2.9.1, 2.9.2, 2.10.0, 2.10.1, 2.10.2, 2.11.0, and 2.11.1. These versions correspond to the following Cloudera Manager versions: 5.10.0, 5.10.1, 5.10.2, 5.11.0, 5.11.1, 5.11.2, 5.12.0, and 5.12.1.

This problem prevents new lineage information from being added to existing and new Hive entities. See the Knowledge Base article [EndPoint2 must not be null or empty](#) for more information.

Cloudera Issue: NAV-5661

Multi-cluster support

Navigator was not able to uniquely identify more than one cluster in installations of Cloudera Manager that supported multiple clusters. This problem only occurs for Cloudera Manager deployments of version 5.12.0, 5.12.1 or 5.13.0. This problem does not affect single cluster deployments or Altus clusters.

This issue has been resolved. The upgrade process to this release of Cloudera Manager resolves the problem.



Note: If your multi-cluster installation is configured for high-availability, you will see exceptions thrown by Navigator while linking HDFS entities. This is a known issue for Cloudera Manager 5.12.0, 5.12.1, 5.13.0 and 5.13.1 releases.

Cloudera Issue: NAV-5615

Navigator did not show metadata content for Isilon

Navigator failed to show file system results in a secure cluster backed by a non-HDFS file system, such as EMC Isilon.

This issue has been resolved.

Cloudera Issue: NAV-5375

Spark lineage not shown in Navigator for HDFS HA

When HDFS is configured for high availability, Spark lineage didn't show in Navigator. In these cases, the log includes index out of bounds exceptions such as the following:

```
com.cloudera.nav.pushextractor.spark.SparkPushExtractor Error extracting Spark operation.
java.lang.ArrayIndexOutOfBoundsException
```

This issue has been resolved.

Cloudera Issue: NAV-4874

Issues Fixed in Cloudera Navigator 2.11.1

Navigator Audit Server build version not available

Navigator Audit Server and Navigator Metadata Server have different build versions and numbers. Navigator Metadata Server build information has always been available in the system logs and on the About page of the Cloudera Navigator console. Build information may be needed by Cloudera support to troubleshoot issues.

This issue has been resolved. With this release, Navigator Audit Server also records its build information in its log files.

Cloudera Issue: NAV-5167

Hue links were missing from Navigator console

The links to Hue from entities in the Navigator console were not generated. Each time the console attempted to show a link, the Navigator Metadata Server log included an error similar to the following:

```
10:05:32.539 AM ERROR HueUtility
[qtp778450792-57]: Unable to get Hue link from CM.
java.lang.LinkageError: loader constraint violation: when resolving method
"com.cloudera.nav.server.NavOptions.getPreferredBaseHueURLForCluster(Ljava/lang/String;)Lcom/google/common/base/Optional;"
the class loader (instance of org/eclipse/jetty/webapp/WebAppClassLoader) of the current
class, com/cloudera/nav/utils/HueUtility, and the class loader (instance of
sun/misc/Launcher$AppClassLoader) for resolved class, com/cloudera/nav/server/NavOptions,
have different Class objects for the type
ter(Ljava/lang/String;)Lcom/google/common/base/Optional; used in the signature{
```

This issue has been resolved.

Cloudera Issue: NAV-4756

Issues Fixed in Cloudera Navigator 2.11.0

Navigator Embedded Solr could reach its limit on number of documents it could store

Navigator Metadata Server extracts HDFS entities by performing a one-time bulk extraction and then switching to incremental extraction. In previous releases, a problem caused HDFS bulk extraction to be run more than one time, resulting in duplicate relations created for HDFS. Over time, embedded Solr runs out of document IDs that it can assign to new relations and fails with following error:

```
Caused by: java.lang.IllegalArgumentException: Too many documents, composite IndexReaders
cannot exceed 2147483519
```

When this happened, Navigator stopped any more extraction of data as no new documents can be added to Solr.

In addition, see [Repairing metadata in the storage directory after upgrading](#).

Fixed in Cloudera Navigator 2.9.2, 2.10.1, 2.11.0

Cloudera Issue: NAV-5600

Solr error message (immense term in queryText field) during upgrade

In some cases, Navigator Metadata Server failed to start after upgrading to Cloudera Navigator 2.10 and displayed messages similar to the following (only relevant details shown):

```
2017-07-01 14:27:27,905 ERROR org.apache.solr.servlet.SolrDispatchFilter
[qtp300983713-60]: null:org.apache.solr.common.SolrException:
Exception writing document id 0c02400d332854fa44a352235b42857c to the index;
. . .
Caused by: java.lang.IllegalArgumentException: Document contains at least one
immense term in field="queryText" (whose UTF8 encoding is longer than the max
```

```
length 32766)
. . .
```

This issue was caused by the underlying Java library that was writing Solr documents (index and searchable entities) during the upgrade. Special characters, such as carriage returns (ASCII 32) were being mishandled and generating more characters than the source element.

This issue has been resolved in this release. The `BinaryRequestWriter`, which was already being used to write relations, is now being used to handle elements.

Cloudera Issue: NAV-4655

Harmless (but abundant) error messages cause frequent log roll-over

An error message was being generated needlessly by Cloudera Navigator when filtering by source. The error could have been ignored but it was being logged and causing logs to roll-over too quickly.

```
2017-06-23 10:58:32,967 ERROR com.cloudera.nav.hdfs.client.InotifyClient
[CDHExecutor-0-CDHUrlClassLoader@5177874]: Error handling event (txid: 1627074917):
Renamed /hbase/data/default/ECDS_PROD1_PARTY_RELATED/ccefb... to
/hbase/archive/data/default/ECDS_PROD1_PARTY_RELATED/cbdka... at time 1493302572987
```

This issue has been resolved in this release. Internally, the logging mechanism has been improved so that it uses a more robust approach to logging, the `ThrottlingLogger` utility class.

Cloudera Issue: NAV-4654

Disk space issue for Navigator Audit Server instances backed by Oracle

Cloudera Navigator deployments using Oracle database as the backing store for Navigator Audit Server were consuming lots of disk space for example, by the `NAV_FILE_INFO_STAGING`. The issue was caused by an unclosed connection and is related to the connection issue discussed in [Navigator Audit Server \(Analytics\) table locking issue \(HDFS activity loader\)](#) on page 2757. However, this issue is specific to Oracle-backed Cloudera Navigator instances.

This issue has been resolved in this release. The connections to Oracle database are now handled properly.

Cloudera Issue: NAV-4467

Audit pipeline test and Oracle error (ORA-24816) issue

Executing queries of a certain size (4000 bytes) in some cases (for example, Impala) would disrupt the Audit Pipeline test thereby registering bad health on the Impalad (Impala daemon) and the Navigator Audit Server logging the following Oracle database error:

```
[1911239784@NavigatorServer-4858]: SQL Error: 24816, SQLState: 99999
[1911239784@NavigatorServer-4858]: ORA-24816: Expanded non LONG bind data
supplied after actual LONG or LOB column...
```

This issue was caused by bind value length potentially greater than 4000 bytes followed by binding for LOB or LONG.

This issue has been resolved in this release. Navigator Audit Server now properly handles the order of the binds needed by Oracle databases.

Cloudera Issue: NAV-4194

Navigator Audit Server (Analytics) table locking issue (HDFS activity loader)

Executing multiple statements sometimes resulted in a mismatch between the connection used to execute the statement and the connection used for the commit. The mismatch would result in the server throwing an error message such as the following:

```
[AnalyticsStagingService]: Could not load HDFS activity data for analytics
org.springframework.dao.CannotAcquireLockException: StatementCallback; SQL
[TRUNCATE TABLE NAV_FILE_INFO_STAGING]...
```

This issue has been resolved in this release. The same connection is now used to execute queries and for commits.

Fixed in Cloudera Navigator 2.8.3, 2.9.2, 2.10.1, 2.11.0

Cloudera Issue: NAV-4026

Invoking auth/group API with invalid role generates NPE and BAD_REQUEST Error

Invoking the `auth/group` endpoint with an invalid role as one of the user-provided parameters causes the server to throw an NPE (null pointer exception), resulting in response from server of HTTP 500 (Internal Server Error).

This issue has been resolved in this release. The server code now checks for this condition (avoiding NPE) and replies with HTTP 400 (BAD_REQUEST) instead.

Cloudera Issue: NAV-3910

Bulk extraction being triggered can result in orphaned relations

In some cases, HDFS bulk extraction was being triggered due to an error in incremental extraction. Specifically, Cloudera Navigator was resetting the long identity of HDFS elements resulting in orphaned relations. This issue has been resolved. A new background task looks for and repairs all orphaned relations.

Fixed in Cloudera Navigator 2.9.2, 2.10.1, 2.11

Cloudera Bug: NAV-4479

Issues Fixed in Cloudera Navigator 2.10.2

Lineage relations were missing for some views with "Endpoint2 must not be null or empty" error

Lineage relations were missing for some views and the Navigator Metadata Server log includes errors with a stack trace that looks similar to the following:

```
2017-10-25 19:00:40,780 ERROR com.cloudera.nav.persist.impl.CompositeLinker
[CDHExecutor-0-CDHUrlClassLoader@381e3]: Internal error while linking.
java.lang.IllegalStateException: EndPoint2 must not be null or empty
```

While Navigator determined lineage relations, it found a relation that was created with the target entity missing, hence "Endpoint2 is null." This situation can occur when a Hive view is altered without changing the columns extracted from the source table, such as when the name or other table metadata on the view is changed. For example, the following operation would produce the error:

```
alter view customers_sales_sw set tblproperties("region"="southwest");
```

This problem can occur in Navigator Metadata Server versions 2.9.0, 2.9.1, 2.9.2, 2.10.0, 2.10.1, 2.10.2, 2.11.0, and 2.11.1. These versions correspond to the following Cloudera Manager versions: 5.10.0, 5.10.1, 5.10.2, 5.11.0, 5.11.1, 5.11.2, 5.12.0, and 5.12.1.

This problem prevents new lineage information from being added to existing and new Hive entities. See the Knowledge Base article [EndPoint2 must not be null or empty](#) for more information.

Cloudera Issue: NAV-5661

Cloudera Navigator Data Management Release Notes

SqoopJobExtractor failures for certain SQL constructs

Certain types of SQL queries that generated sub-operations did not retrieve and display lineage for columns, emitting errors in the Cloudera Navigator console and SQLParserException errors in the Navigator Metadata Server log. For example:

```
...Extractor com.cloudera.nav.sqoop.extractor.SqoopJobExtractor failed for
job_1456104915937_0874
java.lang.RuntimeException: com.akiban.sql.parser.SQLParserException: Encountered " "("
"( "" at line1, column 53.
Was expecting one of: "cs" ...
```

The root cause of this issue is that the open source parser (akiban) cannot parse certain queries that result in sub-operations, specifically those resulting from Sqoop query IMPORT statements. This issue applies to Sqoop1 query IMPORT statements only, not Sqoop EXPORT and TABLE IMPORT statements.

This issue has been resolved in this release. The SQLParserException is caught and the Cloudera Navigator console renders the message shown below rather than breaking the lineage diagram:

	Details	Lineage
--	6	0
Columns	Operation Executions	Inputs
	0	Outputs
▼ Columns		
Navigator was unable to parse the query to determine columns for this entity		
▶ Operation Executions (6)		
▶ Inputs (0)		

Fixed in Cloudera Navigator 2.10.2, 2.9.2

Cloudera Bugs: NAV-4718, NAV-4233, NAV-3142

Issues Fixed in Cloudera Navigator 2.10.1

Navigator Embedded Solr could reach its limit on number of documents it could store

Navigator Metadata Server extracts HDFS entities by performing a one-time bulk extraction and then switching to incremental extraction. In previous releases, a problem caused HDFS bulk extraction to be run more than one time, resulting in duplicate relations created for HDFS. Over time, embedded Solr runs out of document IDs that it can assign to new relations and fails with following error:

```
Caused by: java.lang.IllegalArgumentException: Too many documents, composite IndexReaders
cannot exceed 2147483519
```

When this happened, Navigator stopped any more extraction of data as no new documents can be added to Solr.

In addition, see [Repairing metadata in the storage directory after upgrading](#).

Fixed in Cloudera Navigator 2.9.2, 2.10.1, 2.11.0

Cloudera Issue: NAV-5600**Bulk extraction being triggered can result in orphaned relations**

In some cases, HDFS bulk extraction was being triggered due to an error in incremental extraction. Specifically, Cloudera Navigator was resetting the long identity of HDFS elements resulting in orphaned relations. This issue has been resolved. A new background task looks for and repairs all orphaned relations.

Fixed in Cloudera Navigator 2.9.2, 2.10.1, 2.11

Cloudera Bug: NAV-4479

Issues Fixed in Cloudera Navigator 2.10.0**Navigator cannot extract S3 metadata with some databases**

Using a database other than MySQL, such as PostgreSQL (the default database) for the Navigator Metadata Server generated Java errors in the extracted data for Amazon S3 because of an incorrect string library. The issue has been resolved.

Cloudera Issue: NAV-4342

Metadata added to Amazon S3 files persists in Navigator

Once added to Amazon S3 data and extracted by Navigator, custom metadata (key-value pairs) remain with the data captured by Navigator, even after that metadata may have been deleted from the Amazon S3 files. Navigator continues to show the deleted metadata for that file. This issue has been resolved.

Cloudera Issue: NAV-3928

Managed metadata properties missing on the page

Clicking managed metadata entities in lineage diagrams to display details silently failed, displaying nothing. This issue has been resolved—the condensed details page now displays properly.

Cloudera Issue: NAV-4320

Cluster selection doesn't work properly

In certain cases, cluster facets on the search page did not work correctly.

Cloudera Issue: NAV-4260

Search results retained after deselecting search facet

After searching for a specific type of data using the selector box (the search facet, for example, selecting HDFS), the results populating the browser pane remain despite navigating away from the page or de-selecting the item. This issue has now been resolved. After searching and navigating away from the results or deselecting the data type, resulting display also refreshes, as it should.

Cloudera Issue: NAV-4190

Navigator Audit Server (Analytics) table locking issue (HDFS activity loader)

Executing multiple statements sometimes resulted in a mismatch between the connection used to execute the statement and the connection used for the commit. The mismatch would result in the server throwing an error message such as the following:

```
[AnalyticsStagingService]: Could not load HDFS activity data for analytics
org.springframework.dao.CannotAcquireLockException: StatementCallback; SQL
[TRUNCATE TABLE NAV_FILE_INFO_STAGING]...
```

Cloudera Navigator Data Management Release Notes

This issue has been resolved in this release. The same connection is now used to execute queries and for commits.

Fixed in Cloudera Navigator 2.8.3, 2.9.2, 2.10.1, 2.11.0

Cloudera Issue: NAV-4026

Deleted managed metadata property still visible on Policy page

In certain cases, deleted Managed metadata properties were visible after deletion.

Cloudera Issue: NAV-3778

Issues Fixed in Cloudera Navigator 2.9.2

Navigator Embedded Solr could reach its limit on number of documents it could store

Navigator Metadata Server extracts HDFS entities by performing a one-time bulk extraction and then switching to incremental extraction. In previous releases, a problem caused HDFS bulk extraction to be run more than one time, resulting in duplicate relations created for HDFS. Over time, embedded Solr runs out of document IDs that it can assign to new relations and fails with following error:

```
Caused by: java.lang.IllegalArgumentException: Too many documents, composite IndexReaders cannot exceed 2147483519
```

When this happened, Navigator stopped any more extraction of data as no new documents can be added to Solr.

In addition, see [Repairing metadata in the storage directory after upgrading](#).

Fixed in Cloudera Navigator 2.9.2, 2.10.1, 2.11.0

Cloudera Issue: NAV-5600

Bulk extraction being triggered can result in orphaned relations

In some cases, HDFS bulk extraction was being triggered due to an error in incremental extraction. Specifically, Cloudera Navigator was resetting the long identity of HDFS elements resulting in orphaned relations. This issue has been resolved. A new background task looks for and repairs all orphaned relations.

Fixed in Cloudera Navigator 2.9.2, 2.10.1, 2.11

Cloudera Bug: NAV-4479

Sqoop job extraction failures

Sqoop job extraction was failing because Sqoop table IMPORT operations were not being retrieved from Solr. This issue has been resolved in this release. Sqoop job extraction is now being processed as expected.

Fixed in Cloudera Navigator 2.9.1, 2.9.2

Cloudera Bug: NAV-4225

Navigator Audit Server (Analytics) table locking issue (HDFS activity loader)

Executing multiple statements sometimes resulted in a mismatch between the connection used to execute the statement and the connection used for the commit. The mismatch would result in the server throwing an error message such as the following:

```
[AnalyticsStagingService]: Could not load HDFS activity data for analytics org.springframework.dao.CannotAcquireLockException: StatementCallback; SQL [TRUNCATE TABLE NAV_FILE_INFO_STAGING]...
```

This issue has been resolved in this release. The same connection is now used to execute queries and for commits.

Fixed in Cloudera Navigator 2.8.3, 2.9.2, 2.10.1, 2.11.0

Cloudera Issue: NAV-4026

Issues Fixed in Cloudera Navigator 2.9.1

Sqoop job extraction failures

Sqoop job extraction was failing because Sqoop table IMPORT operations were not being retrieved from Solr. This issue has been resolved in this release. Sqoop job extraction is now being processed as expected.

Fixed in Cloudera Navigator 2.9.1, 2.9.2

Cloudera Bug: NAV-4225

Viewing deleted properties in managed metadata shows incorrect search results

Managed metadata property that was deleted drilled down to incorrect search results. The results are now displayed correctly.

Cloudera Issue: NAV-4116

Edit Managed Metadata link available for users with insufficient privilege

The Edit Managed Metadata link was available to users without sufficient privilege levels to view it (the link requires AUTH_WRITE_MANAGED_METADATA). The link is now available only to users with the correct privileges.

Cloudera Issue: NAV-3985

Some properties that are URLs in the user interface do not function as links

Previously, managed metadata property values and S3 property values that are URLs did not function as links to the URL. These properties are now formatted as clickable links.

Cloudera Issue: NAV-3960

Purge navigation and messages could be confusing

When purging deleted properties in Navigator, the maintenance page message was confusing in some situations, and the Continue link after purge returned you to the general Search page. The purge messages have been clarified and now, after successful purge, the page on which you started redisplay rather than the general Search page.

Cloudera Issue: NAV-3940

Dashboard average values difficult to read

Previously, the dashboard displayed average values in a long format that could be difficult to read. These values have been converted to a format that is more human-readable. For example, the value 33520.2 now reads 33.5k.

Cloudera Issue: NAV-3930

Audit page filter breadcrumbs missing spaces

There were no spaces in the filter breadcrumbs on the Audit page, making the filter difficult to read. Spaces have been added to improve readability.

Cloudera Issue: NAV-3876

Searching for text in the Name filter shows HTML bold tag

When using the Original Name filter in search, characters that you typed in the name box appeared with HTML bold tags in the drop-down results. The extraneous tags no longer appear.

Cloudera Issue: NAV-3765

Cloudera Navigator Data Management Release Notes

Directory contents do not distinguish deleted entities for Amazon S3 entities

In Navigator, search results indicate if an entity has been deleted. However, the Directory Contents information did not indicate that the entity had been deleted; it was indistinguishable from items in Directory Contents that had not been deleted. Directory Contents information now reflects S3 entity deletions.

Cloudera Issue: NAV-3786

Bucket name not included in S3 search results list

In S3 search results, the S3 bucket name was not included with the name of the entity. The S3 bucket name is now displayed.

Cloudera Issue: NAV-3786

Issues Fixed in Cloudera Navigator 2.9.0

Incomplete query strings in Navigator metadata

Storing more than 32K characters in Solr causes an exception, and Navigator previously handled this by truncating any field longer than 32K to prevent the exception.

Now, fields longer than 32K are preserved and no longer truncated.

Cloudera Issues: NAV-3548, NAV-3384

Lineage options and search items are missing

In the Safari browser, lineage options and lineage search could be empty in certain cases.

Cloudera Issue: NAV-3401

Out-of-date information when using Internet Explorer 11

In Internet Explorer 11, cached headers could cause out-of-date information to display.

Issues Fixed in Cloudera Navigator 2.8.3

Audit Event reports exported as CSV randomly contained quote character

When exporting an Audit Events report using the CSV format, a double-quotation mark (") was sometimes being inserted into the query used to collect audit events. For example, in the audit event shown below, the quote turns `_probable` into `_pro"bable` and `_beta` into `_bet"a`:

```
"2016-12-01T12:50:59.966Z",REDACTED,REDACTED,hive,QUERY,
"beta_vat_mart_delta:ds_dup_vrn_pro"bable",true,,,,,,,,,
INSERT INTO exp_vat_oob_load_bet"a.temp_ds_dup_vrn_probable_delta_subset
```

The issue arose when exporting Audit Event reports as CSV format only (not JSON). The issue occurred when the string length of SQL query exceeded 35 characters.

The issue's root cause was an error in processing by the underlying Jackson JSON processor library (jackson-dataformat-csv 2.1.0) used to write the CSV files. The 2.2 version of the library fixes the underlying issue. The issue has been resolved in this release of Cloudera Navigator, which uses the 2.2 Jackson processor library.

Cloudera Issue: NAV-4114

Navigator Audit Server (Analytics) table locking issue (HDFS activity loader)

Executing multiple statements sometimes resulted in a mismatch between the connection used to execute the statement and the connection used for the commit. The mismatch would result in the server throwing an error message such as the following:

```
[AnalyticsStagingService]: Could not load HDFS activity data for analytics  
org.springframework.dao.CannotAcquireLockException: StatementCallback; SQL  
[TRUNCATE TABLE NAV_FILE_INFO_STAGING]...
```

This issue has been resolved in this release. The same connection is now used to execute queries and for commits.

Fixed in Cloudera Navigator 2.8.3, 2.9.2, 2.10.1, 2.11.0

Cloudera Issue: NAV-4026

Issues Fixed in Cloudera Navigator 2.8.1

Some APIs return incorrect error

In some cases, when a user performs an action with managed metadata that is not permitted because of the user role, Navigator issues an Access Denied error message (500) instead of the expected Forbidden error (403).

Cloudera Issue: NAV-3463

Resource path filtering does not working in audit event filter

Lineage options and search items are missing

In the Safari browser, lineage options and lineage search could be empty in certain cases.

Cloudera Issue: NAV-3401

Navigator logout link does not work when SAML is enabled

When SAML is enabled, clicking on the Logout link in Navigator executes only a local browser logout, without terminating the SAML session. As a result, after logging out, a user can log in again from the login page without providing credentials.

Issues Fixed in Cloudera Navigator 2.8.0

Last access time displayed for Hive tables is incorrect

Navigator no longer displays incorrect access time information for Hive.

Cloudera Issue: NAV-99

Deleted items not shown by default when Deleted filter added

Showing deleted items when the Deleted filter was not enabled by default. Deleted items are now shown by default when the Deleted filter is applied.

Cloudera Issue: NAV-441

No message shown when no Groups with Navigator roles are found

A message is now displayed in Navigator when no Groups with Navigator roles are found.

Cloudera Issue: NAV-836

Hive started and ended time fields not shown in results

Previously, started and ended time fields were not displayed when selected. These fields are now displayed in results.

Cloudera Navigator Data Management Release Notes

Cloudera Issue: NAV-1148

Maximum Complexity error message value incorrect

The value in the Maximum Complexity Exceeded error message has been corrected (>3000).

Cloudera Issue: NAV-1246

List of properties in audit filter not sorted

Audit filter properties are now sorted alphabetically.

Cloudera Issue: NAV-1414

Unable to exit Lineage Full Screen mode from Search Results page

Cloudera Issue: NAV-1514

Audit results do not display seconds or fractions seconds

Audit result display granularity has increased to show seconds and fractions of seconds in results.

Cloudera Issue: NAV-1679

Search on lineage page does not work for children

Lineage children are now shown when selected in a lineage search.

Cloudera Issue: NAV-2628

Create property UI does not use class display names

The create property user interface now displays the correct class display name instead of the abbreviated classname.

Cloudera Issue: NAV-2685

Regex entered cannot be confirmed as correct in managed metadata

Navigator now includes regular expression test validation to managed metadata creation.

Cloudera Issue: NAV-2811

Create namespace does not check against reserved names

The Navigator UI now rejects a reserved namespace when creating a new namespace.

Cloudera Issue: NAV-3193

Issues Fixed in Cloudera Navigator 2.7.3

Incomplete query strings in Navigator metadata

Storing more than 32K characters in Solr causes an exception. Navigator now truncates any field longer than 32K to prevent the exception.

Cloudera Issues: NAV-3548, NAV-3384

Issues Fixed in Cloudera Navigator 2.7.2

Long query text alters the display of the search box

Very long values in search strings did not display properly. The issue is fixed in this release.

Cloudera Issue: NAV-3166

Pop-up error in UI in clicking on hive table entity's schema entity

Cloudera Issue: NAV-3579

Lineage Options and Search items are missing

In the Safari browser, lineage options and lineage search could be empty in certain cases.

Cloudera Issue: NAV-3401

Google Analytics didn't load when Navigator is served over HTTPS/SSL

After upgrading to Navigator 2.6.1, Navigator Metadata Server logs included reference to failed loading of Google Analytics. The issue is resolved in this release.

Cloudera Issue: NAV-3200

Issues Fixed in Cloudera Navigator 2.7.1

Oozie and Yarn metadata is not visible for some jobs

The time window for polling for new jobs didn't account for some lag in the system, leaving out some events. The issue is fixed in this release.

Cloudera Issue: NAV-3268

Issues Fixed in Cloudera Navigator 2.7.0

Linker causing high memory consumption

Due to an incorrect internal setting, all results were being returned again in second batch requests to Search (Solr), resulting in high memory usage. The parameter `autoAdjustBatchSize` for the `AbstractSolrManager.query` setting is now correctly set (to `false` rather than `true`) so that subsequent batch requests after an initial request batch the proper number of requests.

Cloudera Issue: NAV-3223

Infinite loop during upgrade possible due to deleted entity query

The `CheckDeletedEntity::start` method was using the `*:*` query to get entities to mark as deleted. This adversely affected performance and made it possible to cause an upgrade to enter an infinite loop if the server crashed or was restarted during an upgrade. This issue has been resolved in this release.

Cloudera Issue: NAV-3111

Managed metadata does not display in lineage diagram

Lineage diagrams display custom metadata, but not managed metadata. This issue has been resolved in this release.

Cloudera Issue: NAV-2919

User-specified relations not properly marked in lineage

User-defined relations were not displaying correctly in the lineage diagrams. This issue has been resolved in this release.

Cloudera Issue: NAV-3137

Exporting lineage endpoints without direction causes errors

Attempting to export lineage endpoints without providing direction (upstream, downstream) in the query returns a server error (error 500) and adds a null pointer exception (NPE) to the log. This issue has been resolved in this release. Empty direction parameter is prevented from being submitted.

Cloudera Navigator Data Management Release Notes

Cloudera Issue: NAV-3155

Placement discrepancy in Cloudera Navigator console for Policies

The Policies **Command Action** section displays above Metadata Assignments for **View** but display below Metadata Assignments for **Edit**. This issue has been resolved in this release.

Cloudera Issue: NAV-3133

Delete VIEW to FILE relations

As part of metadata purge tasks, Navigator failed to delete VIEW to FILE relations. The issue is resolved in this release.

Cloudera Issue: NAV-3106

Logical physical relationships from Hive to YARN should be removed

Upgrading left lineage relationships that incorrectly link Hive entity sources to YARN operation targets. The issue is fixed in this release.

Cloudera Issue: NAV-3105

popover2 binding is positioned incorrectly

The content of the Additional Information pop-over displayed incorrectly in the console pages for creating a new metadata policy. The issue is fixed in this release.

Cloudera Issue: NAV-3071

Show parent when applicable on entity details page

Entities referenced in other entity details pages are now shown with navigation links.

Cloudera Issue: NAV-3034

Entity cannot be deselected in Lineage

Click on a clear section of the lineage canvas to deselect any selected entities or links. Panning or zooming does not deselect the current object.

Cloudera Issue: NAV-3033

Increase zoom out capability in lineage

Controls were added to allow you to zoom further out in the lineage view.

Cloudera Issue: NAV-3032

API/interactive/entities fails when managed metadata properties exist but did not yet have values

Although this problem was not visible in the console, it resulted in a mismatch between what showed in the Navigator console and Navigator stored metadata.

Cloudera Issue: NAV-3012

Entity Details page edit metadata did not show namespace in managed metadata

The namespace name is now available when you hover over the name of the managed metadata property.

Cloudera Issue: NAV-2886

Actions menu should not show Move for deleted entities

The Move menu option is no longer available for deleted entities.

Cloudera Issue: NAV-2727

No Operation/OperationExecution extracted for Hive and HDFS

The previous release introduced a change that caused Navigator communication with services to fail. The issue is fixed in this release.

Cloudera Issue: NAV-2670

numPartitions attribute for Hive should be suppressed

Hive does not produce consistent data for the `tp_numPartitions` property so Navigator no longer shows that value.

Cloudera Issue: NAV-2548

Inconsistent size and block size units

The Search and Analytics pages in the Navigator console used different units for displaying the size of files, tables, and blocks. The issue is fixed in this release.

Cloudera Issue: NAV-2492

Search drop-down width should be the same as the search input field width

The search drop-down now shows more of the previous queries for selection.

Cloudera Issue: NAV-2208

Heap for Navigator Metadata Server increased

The default heap value for Navigator Metadata Server is increased from 696 MB to 1.5 GB.

Cloudera Issue: NAV-2020

Improve labels for Analytics top users and top commands filter

The titles for filters were updated to make them more readable.

Cloudera Issue: NAV-1917

Cross-site Request Forgery (CSRF) vulnerability

Web client requests to Cloudera Navigator were not protected by default from cross-site request forgery attempts. Additional configuration was needed to protect the cluster. That issue has been resolved in this release. CSRF security is now enabled by default.

Cloudera Issue: NAV-1870

Policy earliest start time should be current time

The Navigator console was updated to allow users to set the current time in a policy start time.

Cloudera Issue: NAV-1733

Policy editor line wrapping divides words in the middle

The formatting was changed for in the policy creation interface.

Cloudera Issue: NAV-1702

Date format is different for facets, breadcrumbs, and full query text

Dates and times are now shown consistently using the current locale.

Cloudera Navigator Data Management Release Notes

Cloudera Issue: NAV-1084

Improve rendering for parent elements

This release includes changes in how lineage relationships are shown to improve readability.

Cloudera Issue: NAV-274

"created" date is later than "lastAccessed" date

This release uses the creation time to be the earlier of modified or accessed date when extracting new files.

Cloudera Issue: NAV-39

Added support for managed metadata in the Policy Editor

This release supports policies for managed metadata as well as custom metadata.

Cloudera Issue: NAV-2865

Navigator Metadata Server does not correctly index /hbase/ subpaths

Navigator disabled by default the metadata indexing of /hbase/WALs and /hbase/oldWALs and any entry under /hbase/* entirely.

Previously, /hbase/WALs and /hbase/oldWALs rules were not effective because they were masked by the rule /.*, which was matching everything under /hbase. This has been corrected. Now only nested hidden files and directories under /hbase are discarded.

Cloudera Issue: NAV-3096

Issues Fixed in Cloudera Navigator 2.6.5

Out of memory during file purge

Navigator runs out of memory during purge of files.

Cloudera Issue: NAV-2751

Hive parser issues errors

The Hive parser issued errors during a query in the following cases:

- Multiple struct attributes in a query
- Constants are projected as part of a subquery
- Functions involving constants with an alias

Cloudera Issue: NAV-3402

Incomplete query strings in Navigator metadata

Storing more than 32K characters in Solr causes an exception. Navigator now truncates any field longer than 32K to prevent the exception.

Cloudera Issues: NAV-3548, NAV-3384

Issues Fixed in Cloudera Navigator 2.6.2

Linker causes high memory consumption

When calculating lineage relations, Navigator was pulling more metadata into memory than required. The issue is resolved in this release.

Cloudera Issue: NAV-3223

Navigator Metadata Server fails with nav.batch and a high amount of heap

Cloudera Issue: NAV-3272

java.lang.IllegalArgumentException error causes Navigator to crash

Cloudera Issue: NAV-3244

Class not visible on Managed Metadata review screen

The fields Class and Additional class are not visible on the Review screen when creating or updating a Managed Metadata property.

Cloudera Issue: NAV-3123

Google Analytics didn't load when Navigator is served over HTTPS/SSL

After upgrading to Navigator 2.6.1, Navigator Metadata Server logs included reference to failed loading of Google Analytics. The issue is resolved in this release.

Cloudera Issue: NAV-3200

Column type for HBase qualifier value in Navigator Audit database should be LONGBLOB

The following errors appeared in the Navigator Audit Server logs when HBase audits were propagated to Navigator Audit Server MySQL Database:

```
2016-04-02 03:04:20,783 ERROR org.hibernate.engine.jdbc.spi.SqlExceptionHelper: Data
truncation: Data too long for column 'QUALIFIER' at row 1
2016-04-02 03:04:20,808 WARN com.cloudera.navigator.NavigatorServer: Error persisting
events for service hbase2
```

The type-conversion issue is resolved in this release.

Cloudera Issue: NAV-3147

Issues Fixed in Cloudera Navigator 2.6.1

Lineage returned deleted entities even if the Hide deleted entities option is selected

The issue is fixed in this release.

Cloudera Issue: NAV-2945

Browser sessions could be unexpectedly logged out

When multiple browser tabs are open to the same Metadata Server, an inactive tab can log you out even though you are active in another tab.

Cloudera Issue: NAV-2757

Navigator Metadata Server does not correctly index /hbase/ subpaths

Navigator disabled by default the metadata indexing of /hbase/WALs and /hbase/oldWALs and any entry under /hbase/* entirely.

Previously, /hbase/WALs and /hbase/oldWALs rules were not effective because they were masked by the rule /.*, which was matching everything under /hbase. This has been corrected. Now only nested hidden files and directories under /hbase are discarded.

Cloudera Issue: NAV-3096

Cloudera Navigator Data Management Release Notes

Issues with additional Group facet values in the Search page

If a new value is added to the subfacet SourceType:HDFS->Group, the 'Clear all' actions selects all the values in the Group facet, although they are not included in the query.

The input field expands outside the facet panel.

Adding a new (not suggested) value to the Group facet in Firefox leaves the suggestion drop-down open.

Cloudera Issue: NAV-3017

Logical Physical relationships unexpectedly shown in lineage

HDFS directories are displayed in Hive lineage.

Cloudera Issue: NAV-2984

Lineage rendering sometimes throws NPE when latest operation is selected

Missing start time values in the metadata used to generate lineage relationships causes errors in the Navigator console lineage view. The issue is fixed in this release.

Cloudera Issue: NAV-2969

ClassCastException thrown when a custom entity is encountered

The Navigator API allowed custom entities to be created without some required metadata; the Navigator console produced errors when attempting to display such an entity.

Cloudera Issue: NAV-2946

Mismatch in number of files between Search and Search from Analytics due to deleted files

Search queries invoked from the Analytics page now show the correct number of files.

Cloudera Issue: NAV-2854

On load, the originating entity might not be visible in lineage

If the lineage graph is tall or wide the originating entity might not be initially visible.

Cloudera Issue: NAV-2798

Browser support warning was not working

Internet Explorer 11 was not in the list of supported browsers.

Cloudera Issue: NAV-2780

Lineage could fail or display partial results when related to Workflows entities

Oozie workflows can have multiple parents however lineage cannot display this use case. Workflows are no longer returned for an operation. To get workflow detail, you must launch lineage on the workflow itself.

Cloudera Issue: NAV-2775

Metadata UI could be unreliable when viewing entities with no name

If an entity has no value in the name or originalName properties, they are now set to "(no name)".

Cloudera Issue: NAV-2746

User logging would in some cases be presented with the login screen again

Navigator console logout timing didn't match the server's timing and caused this loop of logout screens. The issue is fixed in this release.

Cloudera Issue: NAV-2731

Lineage load error banner could show on subsequent lineages

Navigator displays two warnings when lineage load fails. The yellow banner displays until dismissed or the page is reloaded even if it does not apply to the current lineage.

Cloudera Issue: NAV-2697

Missing managed metadata property details

Regular expressions and other details are not listed in the detailed managed metadata property view.

Cloudera Issue: NAV-2663

Alignment of logical physical relations does not follow data flow

Logical physical relations should be shown such that they follow data flow. For example, if downstream you hit a logical entity first then it should be logical -> physical. If you hit physical first it should be physical -> logical.

Cloudera Issue: NAV-2639

Input validation could in certain cases show errors after the error was fixed

Native input validation messages on properties in the Edit Metadata dialog keep displaying, for example when working with regular expression patterns.

Cloudera Issue: NAV-2523

Combo box display issues in Internet Explorer

The combo box scroll bar in search facets does not work in Internet Explorer. Clicking it closes the drop-down.

Cloudera Issue: NAV-2522

Hue link missing

When the Cluster Name facet is used, the Hue link is missing in search results.

Cloudera Issue: NAV-1927

Issues Fixed in Cloudera Navigator 2.6.0

Lineage performance suffers when more than 10000 relations are extracted

If more than 10000 relations must be traversed for a lineage diagram, performance suffers. This can occur when there are thousands of files in a directory or hundreds of columns in a table.

Cloudera Issue: NAV-390

In 2.3, the Navigator Search UI is missing the User Defined Property filter

The Navigator Search UI has changed between 2.2 and 2.3 with respect to the User Defined Property filter. In 2.3 it is no longer present.

Workaround: In the Search box type: `up_propertyName:value`. For example, `up_myName:myValue`.

Cloudera Issue: NAV-1691

Displaying the lineage of an entity may consume too much memory and the Metadata Server may run out of memory

Memory leaks have been fixed. However, the Navigator Metadata Server embeds Solr server inside its process, and its memory requirement is proportional to the data that is being indexed. For recommendations on how to configure memory, see [Memory Sizing Considerations for Navigator Metadata Server](#).

Cloudera Issue: NAV-1832

Cloudera Navigator Data Management Release Notes

After upgrade to 2.4.x, only entries related to deleted files are visible

In upgrading, Navigator failed to properly convert existing Hive entities to the new system. The problem involved how Hive query contents were used for creating entity names. The issue is fixed in this release.

Cloudera Issue: NAV-2542

Issues Fixed in Cloudera Navigator 2.5.0

After upgrade to 2.4.x, only entries related to deleted files are visible

In upgrading, Navigator failed to properly convert existing Hive entities to the new system. The problem involved how Hive query contents were used for creating entity names. The issue is fixed in this release.

Cloudera Issue: NAV-2542

Issues Fixed in Cloudera Navigator 2.4.4

Lineage does not render on MacOS using Chrome 48.0.2564.109

This change anticipates Google's deprecation and removal of the API that Navigator used to render lineage in Chrome browsers.

Cloudera Issue: NAV-2459

After upgrade to 2.4.x, only entries related to deleted files are visible

In upgrading, Navigator failed to properly convert existing Hive entities to the new system. The problem involved how Hive query contents were used for creating entity names. The issue is fixed in this release.

Cloudera Issue: NAV-2542

Issues Fixed in Cloudera Navigator 2.4.2



Note: If you are upgrading from a previous version of Cloudera Navigator, Cloudera recommends that you upgrade to version 2.4.3 or higher.

Searching for strings that start with forward slash ignored sometimes

Searching for a string that starts with a forward slash (/) in conjunction with another filter/facet or followed by AND and another search term will result in the string being ignored completely.

Workaround: Enclose the string starting with forward slash in double quotes, for example `"/some/path" AND sourceType:hdfs`

Cloudera Issue: NAV-2022

After upgrade, entities are incorrectly marked as deleted

The issue is resolved. However, upgrading from Cloudera Navigator 2.4.0 or 2.4.1 to 2.4.2 requires manual steps after the upgrade, in addition to the process detailed in the upgrade documentation. Specifically, the HDFS extractor state file (on the host system for each HDFS role instance that Cloudera Navigator extracts from) must be modified as detailed below:

1. Edit the HDFS extractor state file under the Cloudera Navigator storage directory. The default path is:

```
/var/lib/cloudera-scm-navigator/extractorState
```


The HDFS extractor state file contains a JSON object with an attribute `upgrades`, which has a list of upgrades and their statuses. For example:

```
{ "nextTxId":11394669, "lastKnownTransactionTime":1453931555151, "layoutVersion":-60,
"namespaceId":1613692955, "nextExtractionRunId": "df61cf41183dec8947078c228ec505e6##7453",
"upgrades": [
{"taskType":"SET_MISSING_ATTRIBUTES", "status":"SUCCEEDED", "startTime":1453484370390, "endTime":1453484370410, "attempts":1},
{"taskType":"UNDELETE", "status":"SUCCEEDED", "startTime":1453484370391, "endTime":1453484370418, "attempts":1},
{"taskType":"BLOCK_SIZE", "status":"SUCCEEDED", "startTime":1453484370391, "endTime":1453484370427, "attempts":1},
{"taskType":"REPLICATION_COUNT", "status":"SUCCEEDED", "startTime":1453484370391, "endTime":1453484370434, "attempts":1}]}
```

- Find entries in the `upgrades` list for `"taskType": "UNDELETE"`, `"status": "SUCCEEDED"` and change status to failed. For example, this entry:

```
{ "taskType": "UNDELETE", "status": "SUCCEEDED", "startTime": 1453484370391, "endTime": 1453484370418, "attempts": 1 }
```

should be changed to:

```
{ "taskType": "UNDELETE", "status": "FAILED", "startTime": 1453484370391, "endTime": 1453484370418, "attempts": 1 }
```

- Repeat this process for each HDFS extractor state file as needed for your Cloudera Navigator instance.

Cloudera Issue: NAV-2462

On a newly installed cluster, Hive lineage does not work

When Cloudera Manager starts, it does not automatically create the directory `/var/log/hive/lineage`. Because the Agent cannot detect this directory, it fails to set a watch on it. HiveServer2 creates the directory, but Hive metadata extracted to the directory is never read by Agent and sent to Navigator Metadata Server.

Workaround: Do one of the following:

- Use the Cloudera Manager Admin Console to restart HiveServer2 (results in downtime).
- To avoid HiveServer2 downtime, restart the Cloudera Manager Agent on the host running HiveServer2 role instance as follows:

```
sudo service cloudera-scm-agent restart
```

Cloudera Issue: NAV-2111

Entity landing page schema overflow hides name

If the schema type is very long, it can hide the name of the schema entity.

Cloudera Issue: NAV-2083

When Navigator uses the Cloudera Manager authenticator and Cloudera Manager is TLS enabled, authentication fails

When Navigator is configured to use the Cloudera Manager DB for authentication, and Cloudera Manager is TLS enabled, when a user attempts authentication, certificate validation fails with:

```
Caused by: sun.security.validator.ValidatorException: PKIX path building failed:
sun.security.provider.certpath.SunCertPathBuilderException: unable to find valid
certification path to requested target
```

Cloudera Issue: NAV-2178

Analytics page of Cloudera Navigator console includes deleted files

The Cloudera Navigator console **Analytics** page was including deleted files in its counts. This issue has been resolved.

Cloudera Navigator Data Management Release Notes

Cloudera Issue: NAV-2182

Policies created from the Analytics page were not working

Policies created from analytics UI were not working because it filtered on clusterName. Now it filters on sourceId.

Cloudera Issue: NAV-2183

No indication when Entity Details page is not completely loaded

Technical metadata displays immediately but it was not always obvious that loading Entity Details page was still loading. A loading spinner now displays on slow loads.

Cloudera Issue: NAV-2286

Tables might not load on Entity Details page

Tables might not load, depending on whether the schema finished loading before or after tables finish loading.

Cloudera Issue: NAV-2288

Issues Fixed in Cloudera Navigator 2.4.1

Potential Security Vulnerability Resolved

This release resolves a potential security vulnerability that may enable an attacker to execute arbitrary code from a remote machine without requiring authentication. Upgrading to the this release resolves the issue. See the Cloudera Security Bulletin, [Apache Commons Collections deserialization vulnerability](#), for complete details.

Cloudera Issue: NAV-2230

Issues Fixed in Cloudera Navigator 2.4.0

After HDFS upgrade, some changes to HDFS entities may not appear in Navigator

After an HDFS upgrade, Navigator might not detect changes to HDFS entities, such as move, rename, and delete operations, that were recorded only in the HDFS edit logs before the upgrade. This may cause an inconsistent view of HDFS entities between Navigator and HDFS.

Workaround: None.

Cloudera Issue: NAV-730

After upgrading TLS/SSL-enabled Navigator to 5.4.0, Navigator starts in non-TLS/SSL

After upgrading an TLS/SSL-enabled Navigator to 5.4.x, Navigator will start in non-TLS/SSL mode. Before the 5.4.x release of Navigator, TLS/SSL for Navigator was turned on using an advanced configuration snippet. In release 5.4.x, this option is exposed in the Cloudera Manager configuration UI. After upgrade to 5.4.x, this setting overrides the setting in the advanced configuration snippet, disabling TLS/SSL.

Workaround: Remove the setting from the advanced configuration snippet and specify it in the configuration UI directly.

Cloudera Issue: NAV-1536

Navigator login using Cloudera Manager credentials does not work if Cloudera Manager has TLS enabled

With Navigator configured to use Cloudera Manager for authentication and Cloudera Manager using TLS, when a user attempts authentication, certificate validation fails with:

```
Caused by: sun.security.validator.ValidatorException: PKIX path building failed:  
sun.security.provider.certpath.SunCertPathBuilderException: unable to find valid  
certification path to requested target
```

The issue is fixed in this release.

Cloudera Issue: NAV-2207

Memory leak when creating Impala lineage for very big queries (more than 32K characters)

Cloudera Issue: NAV-1832

Disabling auditing in Hive causes exception in HiveServer2

If you disable auditing for Hive, but do not restart the Hive service, HiveServer2 logs the following exception:

```
Hive Internal Error: com.cloudera.navigator.shaded.avro.AvroRuntimeException(Field  
serviceName type:STRING pos:0 does not accept null values)
```

Cloudera Issue: NAV-1316

Lineage is not collected for Hive on Spark jobs

Cloudera Issue: NAV-1309

Handle HDFS root as special case when generating Hue link

Cloudera Issue: NAV-1817

Entity description does not allow multiple lines

Cloudera Issue: NAV-422

Expanding entity causes unintuitive lineage layout

Cloudera Issue: NAV-1438

Selected Pig field not shown in lineage

When you view lineage for a Hive field, the Hive parent is expanded so you see the field. However, if there is a field of a Pig table inside a Pig operation execution, only the operation execution is expanded, not the table. The field is not visible and it is not obvious to which entity it belongs.

Cloudera Issue: NAV-1440

Pig inside Oozie not shown in lineage view

Cloudera Issue: NAV-1662

Page headers missing after selecting "Show Full Page" in lineage diagram

Cloudera Issue: NAV-1681

Facet does not visually retain selection after refresh

After reloading a page, some selections are not rendered as selected even though the search string itself is still correct.

Cloudera Issue: NAV-1913

Link and parent Hive table not displayed

If you launch a job with two input directories that are also both Hive tables, the Hive table is shown for one of the directories but not the other. If you expand the other directory, the Hive table is shown.

Cloudera Issue: NAV-1794

Links disappear when expanding Pig job run by Oozie

Cloudera Issue: NAV-1819

Cloudera Navigator Data Management Release Notes

Expanding a directory folder obscures the Hive table it is linked to

Cloudera Issue: NAV-1908

Source filter missing from filter list

Cloudera Issue: NAV-1159

Subfacets should be enabled and disabled based on values in facet type

Cloudera Issue: NAV-1051

Tags and key-value pairs should be sorted in the policy editor

Cloudera Issue: NAV-1760

Create Policy link should not be available for a policy viewer role

In this release, the Create Policy button does not appear for users who do not have a role that provides permission to create or update policies.

Cloudera Issue: NAV-1972

Impala auditing does not support filtering during event capture

Impala auditing does not support filtering during event capture.

Severity: Low

Workaround: None.

Cloudera Issue: OPSAPS-15372

Issues Fixed in Cloudera Navigator 2.3.9

Apache Commons Collections Deserialization Vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera (“Cloudera Products”), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled “Arbitrary remote code execution with InvokerTransformer” and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower

Users affected: All

Severity (Low/Medium/High): High

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Navigator 2.3.9.

Issues Fixed in Cloudera Navigator 2.3.8

Disabling auditing in Hive causes exception in HiveServer2

If you disable auditing for Hive, but do not restart the Hive service, HiveServer2 logs the following exception:

```
Hive Internal Error: com.cloudera.navigator.shaded.avro.AvroRuntimeException(Field  
serviceName type:STRING pos:0 does not accept null values)
```

Cloudera Issue: NAV-1316

NPE during YARN extraction

The extractor no longer throws a NPE during YARN extraction.

Cloudera Issue: NAV-2016

Issues Fixed in Cloudera Navigator 2.3.3

Upgrade to 2.3.1 does not create column "NAME" in NAVMS_AUDIT_EVENTS

The upgrade process neglected to include a script to create the "NAME" column.

Cloudera Issue: NAV-1541

Navigator Search UI hangs due to blocked log4j thread

The Metadata Server log4j logger blocked due to a defect in how old logs are deleted. The logger has been replaced and the Cloudera Manager Agent deletes old log files.

Cloudera Issue: NAV-1539

Expanding Pig lineage can cause overlap

When you expand a Pig lineage diagram, relationship lines can sometimes obstruct one another.

Cloudera Issue: NAV-1545

Audit reports subject to cross-site scripting

Navigator does not do input validation on report names, which allows users to enter arbitrary code in this field. The next time a report name is viewed, the script is executed.

Cloudera Issue: NAV-1502

Session timeout results in a redirect using HTTP basic authentication

Upon session timeout, Navigator redirects to HTTP basic authentication instead of redirecting to the main login page.

Cloudera Issue: NAV-1506

"POODLE" vulnerability on TLS/SSL enabled ports

The POODLE (Padding Oracle On Downgraded Legacy Encryption) attack takes advantage of a cryptographic flaw in the obsolete SSLv3 protocol, after first forcing the use of that protocol. The only solution is to disable SSLv3 entirely. This requires changes across a wide variety of components of Cloudera Navigator in 2.3.x and all earlier versions. Cloudera Navigator 2.3.3 provides these changes for Cloudera Navigator 2.3.x deployments. All Cloudera Navigator users should upgrade to 2.3.3 as soon as possible. For more information, see the [Cloudera Security Bulletin](#).

Cloudera Issue: NAV-1501

LDAP group searches appear slow because of web UI behavior

Searching for LDAP groups in the web UI can appear slow because every time a character changes in the search field, an onchange Javascript event triggers a new LDAP search.

Cloudera Navigator Data Management Release Notes

Cloudera Issue: NAV-1491

Solr audit fields are incorrectly shown with capital letters in the web UI

Solr audit fields such as **sub_operation** and **entity_id** are rendered with capital letters and spaces such as **Sub Operation** or **Entity ID**.

Cloudera Issue: NAV-1481

Issues Fixed in Cloudera Navigator 2.3.1

Link disappears when expanded

Sometimes when you expand a lineage diagram, a link disappears.

Cloudera Issue: NAV-1416

Operations on canary files should be filtered out

Cloudera Issue: NAV-1419

Hive views have same icon as Hive fields in lineage

Cloudera Issue: NAV-1365

Facet count should show (0) instead of (-) when there are no matching entities

Facet counts show (-) when the Metadata Server does not return a value. It should show (0).

Workaround: None.

Cloudera Issue: NAV-1271

Column lineage does not display

Column lineage does not display when its parent is automatically expanded.

Workaround: Redisplay the lineage by clicking the parent entity.

Cloudera Issue: NAV-1380

Kite dataset extraction fails when HDFS HA is enabled

Workaround: None.

Cloudera Issue: NAV-1383

Exporting audit reports to CSV does not work

You can export audits to JSON with a limit of 10,000 audits.

Cloudera Issue: NAV-1332

Sentry auditing does not work if the Python version is lower than 2.5

Cloudera Issue: NAV-1080

Sqoop sub-operations don't display in schema view of lineage

Workaround: None.

Cloudera Issue: NAV-1211

Issues Fixed in Cloudera Navigator 2.3.0

There were no issues fixed in Cloudera Navigator 2.3.0.

Issues Fixed in Cloudera Navigator 2.2.9

Apache Commons Collections Deserialization Vulnerability

Cloudera has learned of a potential security vulnerability in a third-party library called the [Apache Commons Collections](#). This library is used in products distributed and supported by Cloudera (“Cloudera Products”), including core Apache Hadoop. The Apache Commons Collections library is also in widespread use beyond the Hadoop ecosystem. At this time, no specific attack vector for this vulnerability has been identified as present in Cloudera Products.

In an abundance of caution, we are currently in the process of incorporating a version of the Apache Commons Collections library with a fix into the Cloudera Products. In most cases, this will require coordination with the projects in the Apache community. One example of this is tracked by [HADOOP-12577](#).

The Apache Commons Collections potential security vulnerability is titled “Arbitrary remote code execution with InvokerTransformer” and is tracked by [COLLECTIONS-580](#). MITRE has not issued a CVE, but related [CVE-2015-4852](#) has been filed for the vulnerability. CERT has issued [Vulnerability Note #576313](#) for this issue.

Releases affected: CDH 5.5.0, CDH 5.4.8 and lower, Cloudera Manager 5.5.0, Cloudera Manager 5.4.8 and lower, Cloudera Navigator 2.4.0, Cloudera Navigator 2.3.8 and lower, Cloudera Navigator 2.2.8 and lower.

Users affected: All

Severity (Low/Medium/High): High

Impact: This potential vulnerability may enable an attacker to execute arbitrary code from a remote machine without requiring authentication.

Immediate action required: Upgrade to Cloudera Navigator 2.2.9.

Issues Fixed in Cloudera Navigator 2.2.4

“POODLE” vulnerability on TLS/SSL enabled ports

The POODLE (Padding Oracle On Downgraded Legacy Encryption) attack takes advantage of a cryptographic flaw in the obsolete SSLv3 protocol, after first forcing the use of that protocol. The only solution is to disable SSLv3 entirely. This requires changes across a wide variety of components of Cloudera Navigator in 2.2.x and all earlier versions. Cloudera Navigator 2.2.4 provides these changes for Cloudera Navigator 2.2.x deployments. All Cloudera Navigator users should upgrade to 2.2.4 as soon as possible. For more information, see the [Cloudera Security Bulletin](#).

Cloudera Issue: NAV-1501

Issues Fixed in Cloudera Navigator 2.2.3

Navigator Audit Server reports invalid null characters in HBase audit events when using the PostgreSQL database

Navigator Audit Server reports invalid null characters in HBase audit events when using the PostgreSQL database. HBase allows null characters in qualifiers, so now Navigator escapes them.

Cloudera Issue: NAV-1286

Oozie extractor throws too many Boolean clauses exception

Cloudera Issue: NAV-1203

Issues Fixed in Cloudera Navigator 2.2.2

The audit reports UI now returns results when there are a large number of audit records

The audit reports UI was not returning results when there were a large number of audit records matching a particular time period, especially when the period included multiple days. The UI is now also much more responsive.

Cloudera Issue: NAV-888

Issues Fixed in Cloudera Navigator 2.2.1

Browser autocomplete no longer enabled before authentication

Form fields before authentication in the application have auto-complete enabled. Any user using the same computer would be able to see information entered by a previous user.

Cloudera Issue: NAV-924

Navigator Web UI no longer exposes paths to directory listing/forceful browsing

The web server is configured to display the list of files contained in this directory. This is not recommended because the directory may contain files that are not normally exposed through links on the web site.

Cloudera Issue: NAV-923

Navigator Audit Server no longer throws OOM for very long Impala queries

Cloudera Issue: NAV-952

Issues Fixed in Cloudera Navigator 2.2.0

If Hue is added after Navigator, search results do not have links to Hue

In the Metadata UI, search results contain links to an appropriate application in Hue. However, if you add a Hue service after Navigator roles, there will be no links to Hue.

Workaround:

1. Set the cluster's display name and name properties to be the same:
 - a. Get cluster's name and display name using following API: `http://hostname:7180/api/v6/clusters`.
 - b. In the Cloudera Manager Admin Console, at the right of the cluster name, click the down arrow and select **Rename Cluster**. Set the cluster display name to match its name.
2. Restart the Navigator Metadata server.

Cloudera Issue: NAV-745

Issues Fixed in Cloudera Navigator 2.1.6

When auditing is enabled, the Cloudera Manager Agent may become slow or get stuck when responding to commands and when sending heartbeats to Cloudera Manager Server

This issue can occur when Cloudera Navigator auditing is turned on. The auditing code reads audit logs and sends them to the Audit Server. It acquires a lock to protect the list of roles being audited. The same list is also modified by the Cloudera Manager Agent's main thread when a role is started or stopped. If the Audit thread takes too much time to send audits to the Audit Server (which can happen if there is backlog of audit logs), it starves the main Agent thread. This causes the main Agent thread to not send heartbeats and to not respond to commands from the Cloudera Manager Server.

Issues Fixed in Cloudera Navigator 2.1.5

Navigator Audit Server reports invalid null characters in HBase audit events when using the PostgreSQL database

Navigator Audit Server reports invalid null characters in HBase audit events when using the PostgreSQL database. HBase allows null characters in qualifiers, so now Navigator escapes them.

Cloudera Issue: NAV-1286

Oozie extractor throws too many Boolean clauses exception

Cloudera Issue: NAV-1203

Issues Fixed in Cloudera Navigator 2.1.4

The audit reports UI now returns results when there are a large number of audit records

The audit reports UI was not returning results when there were a large number of audit records matching a particular time period, especially when the period included multiple days. The UI is now also much more responsive.

Cloudera Issue: NAV-888

Issues Fixed in Cloudera Navigator 2.1.2

Search results in Navigator now have links to Hue

In the Metadata UI, search results contain links to an appropriate application in Hue. The links may be missing for either of the following reasons:

- Hue was added after Navigator Metadata server was started.
- The cluster name and display name are different.

Workaround:

1. Set the cluster's display name and name properties to be the same:
 - a. Get cluster's name and display name using following API: `http://hostname:7180/api/v6/clusters`.
 - b. In the Cloudera Manager Admin Console, at the right of the cluster name, click the down arrow and select **Rename Cluster**. Set the cluster display name to match its name.
2. Restart the Navigator Metadata server.

Cloudera Issues: NAV-745, NAV-866

Browser autocomplete no longer enabled before authentication

Form fields before authentication in the application have auto-complete enabled. Any user using the same computer would be able to see information entered by a previous user.

Cloudera Issue: NAV-924

Navigator Web UI no longer exposes paths to directory listing/forceful browsing

The web server is configured to display the list of files contained in this directory. This is not recommended because the directory may contain files that are not normally exposed through links on the web site.

Cloudera Issue: NAV-923

Navigator Audit Server no longer throws OOM for very long Impala queries

Cloudera Issue: NAV-952

Issues Fixed in Cloudera Navigator 2.1.1

LDAP lookups in Active Directory to resolve group membership are now working

Cloudera Issue: NAV-784

Dropping a Hive table and creating a view with same name or vice versa no longer raises an error

Cloudera Issue: NAV-759

HDFS extraction now works after upgrading CDH from 5.1 to 5.2

Cloudera Issue: NAV-709

Setting a property in the Hue advanced configuration snippet no longer throws a "too many Boolean clauses" error in Navigator Metadata

Cloudera Issue: NAV-549

Issues Fixed in Cloudera Navigator 2.0.5

Memory leak in Navigator Audit Server due to error during batch operations

Cloudera Issue: NAV-984

Issues Fixed in Cloudera Navigator 2.0.3

Dropping a Hive table and creating a view with same name or vice versa no longer raises an error

Cloudera Issue: NAV-759

Setting a property in the Hue advanced configuration snippet no longer throws a "too many Boolean clauses" error in Navigator Metadata

Cloudera Issue: NAV-549

Issues Fixed in Cloudera Navigator 2.0.2

HBase auditing initialization failure can prevent region opening indefinitely

Cloudera Issue: NAV-512

Issues Fixed in Cloudera Navigator 2.0.1

Hive View to Table lineage is missing

The lineage of the underlying tables does not appear in the lineage view.

Workaround: Launch lineage on the underlying tables directly.

Cloudera Issue: NAV-432

The "allowed" query selector is missing from the audit REST API

Queries such as `http://hostname:7180/api/v7/audits?maxResults=10&query=allowed==false` are now supported.

Workaround: None.

Cloudera Issue: NAV-404

Metadata in metadata files is not processed

Workaround: None.

Cloudera Issue: NAV-490

Lineage does not work when launched on a field

When you launch lineage on a field (Hive column or Pig field, or a Sqoop sub-operation), the UI displays the initial graph properly. However if you expand the parent item (Hive table in case of a Hive column), then things start disappearing from the lineage diagram.

Workaround: None.

Cloudera Issue: NAV-365

When you specify an end date for the `created` property, no results are returned.

Workaround: Clear the end date control or specify an end date of `TO+*%5D%22%7D`.

Cloudera Issue: NAV-335

Navigating back to the parent entity in a Pig lineage diagram sometimes displays the error: Cannot read property 'x' of undefined.

Workaround: None.

Cloudera Issue: NAV-298

Issues Fixed in Cloudera Navigator 2.0.0

The last accessed time for Hive table is incorrect.

Workaround: None.

Cloudera Issue: OPSAPS-17227

Pig job that has relations with self is unreadable in lineage view.

The Metadata UI currently does not handle situation where there is a data-flow relation between elements that are also related via parent-child relation.

Workaround: None.

Cloudera Issue: OPSAPS-18836

If auditing is enabled, during an upgrade from Cloudera Manager 4.6.3 to 4.7, the Impala service won't start.

Impala auditing requires the Audit Log Directory property to be set, but the upgrade process fails to set the property.

Workaround: Do one of the following:

- Stop the Cloudera Navigator Audit server.
- Ensure that you have a Cloudera Navigator license and manually set the property to `/var/log/impalad/audit`.

Cloudera Issue: OPSAPS-15333

Empty box appears over the list of results after adding a tag to a file

When tags are added to an entity, in some cases a white box remains after pressing **Enter**.

Workaround: Refresh the page to remove the artifact.

Cloudera Issue: NAV-87

Issues Fixed in Cloudera Navigator 1.2.0

Certain complex multi-level lineages, such as directory/file and database/table, may not be fully represented visually.

Workaround: None.

Cloudera Issue: OPSAPS-18207

Cloudera Navigator Key Trustee Server Release Notes

These release notes provide information on the new and changed features, known issues, and fixed issues for Cloudera Navigator Key Trustee Server.

New Features and Changes in Cloudera Navigator Key Trustee Server

The following sections describe what's new and changed in each Cloudera Navigator Key Trustee Server release.



Note: Releases 5.4.3 and 5.4.9 are the only maintenance releases of Key Trustee Server 5.4.

Release 5.5.2 is the only maintenance release of Key Trustee Server 5.5.

There is no Release 5.6 of Key Trustee Server.

There are no maintenance releases of Key Trustee Server 5.7.

There are no maintenance releases of Key Trustee Server 5.8.

There are no maintenance releases of Key Trustee Server 5.9.

What's New in Cloudera Navigator Key Trustee Server

The following sections describe what's new in each Cloudera Navigator Key Trustee Server release.

What's New in Cloudera Navigator Key Trustee Server 5.15.0

An issue has been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.15.0](#) on page 2790.

What's New in Cloudera Navigator Key Trustee Server 5.14.0

An issue has been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.14.0](#) on page 2791.

What's New in Cloudera Navigator Key Trustee Server 5.13.0

An issue has been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.13.0](#) on page 2791.

What's New in Cloudera Navigator Key Trustee Server 5.12.0

- Key Trustee Server supports RHEL 6.9.
- Key Trustee Server supports CentOS 6.9.

What's New in Cloudera Navigator Key Trustee Server 5.11.0

- Key Trustee Server supports RHEL 7.4.
- Key Trustee Server supports OEL 7.4.

An issue has been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.11.0](#) on page 2791.

What's New in Cloudera Navigator Key Trustee Server 5.10.1

An issue has been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.10.1](#) on page 2791.

What's New in Cloudera Navigator Key Trustee Server 5.10.0

- Key Trustee Server supports RHEL 7.3.
- Key Trustee Server supports OEL 6.8, 7.2, and 7.3.
- Key Trustee Server supports modifying the supported ciphers and protocols for TLS.

Cloudera Navigator Key Trustee Server Release Notes

What's New in Cloudera Navigator Key Trustee Server 5.9.0

- Key Trustee Server supports RHEL 6.8.
- Running the [Key Trustee Server backup script](#) now backs up necessary hardware security module (HSM) configuration files.
- Key Trustee Server supports rolling restart in Cloudera Manager.

What's New in Cloudera Navigator Key Trustee Server 5.8.0

- When adding the parcel-based Key Trustee Server service for the first time, Cloudera Manager automatically backs up Key Trustee Server locally and schedules ongoing hourly local backups using `cron`.
- The Key Trustee Server backup script (`ktbackup.sh`) adds a new option, `--roll`, which specifies the number of backups to retain.

An issue has also been fixed. For more information, see [Issues Fixed in Cloudera Navigator Key Trustee Server 5.8.0](#) on page 2791.

What's New in Cloudera Navigator Key Trustee Server 5.7.0

- A backup script (`ktbackup.sh`) is included with Key Trustee Server.
- Parcel-based Key Trustee Server logs an error message at startup if the `keytrustee.conf` file is malformed.
- Error logging is improved for connection and certificate errors with Key Trustee Server connecting to Key HSM.

A number of issues have also been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.7.0](#) on page 2792.

What's New in Cloudera Navigator Key Trustee Server 5.5.2

- The `ktadmin` command has a new `--passphrase` option to allow migration of existing keys from a Key Trustee Server with a password-protected private key to an HSM.

An issue has also been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.5.2](#) on page 2792.

What's New in Cloudera Navigator Key Trustee Server 5.5.0

- Key Trustee Server supports RHEL 7.
- Key Trustee Server supports password-protected SSL certificates.
- TLS/SSL certificate file paths are configurable in Cloudera Manager.

A number of issues have also been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.5.0](#) on page 2792.

What's New in Cloudera Navigator Key Trustee Server 5.4.9

An issue has been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.4.9](#) on page 2794.

What's New in Cloudera Navigator Key Trustee Server 5.4.3

A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator Key Trustee Server 5.4.3](#) on page 2795.

What's New in Cloudera Navigator Key Trustee Server 5.4.0

- Key Trustee Server can be installed and managed using Cloudera Manager.
- Existing keys stored in Key Trustee Server can be migrated to a Hardware Security Module (HSM) using Navigator Key HSM.

Changed Features and Behaviors in Cloudera Navigator Key Trustee Server

The following sections describe what's changed in each Cloudera Navigator Key Trustee Server release.

What's Changed in Cloudera Navigator Key Trustee Server 5.8.0

- Attempting to start the `keytrusteed` service using the `service` command when Key Trustee Server is already running now reports success with a message indicating that the service is already running.

What's Changed in Cloudera Navigator Key Trustee Server 5.7.0

The `ktadmin keyhsm` command can be used to update Key Trustee Server certificate information for Key HSM

The `ktadmin keyhsm` command can be used to update Key Trustee Server certificate information for Key HSM instead of manually modifying `keytrustee.conf`.

What's Changed in Cloudera Navigator Key Trustee Server 5.5.0

Key Trustee Server can be installed using Cloudera Manager without a CSD

Key Trustee Server no longer requires a Custom Service Descriptor (CSD) file.

What's Changed in Cloudera Navigator Key Trustee Server 5.4.0

All processes run as a single user

All Key Trustee Server processes now run as a single non-root user (`keytrustee` by default).

Key Trustee Server processes use a single port

Key Trustee Server processes (with the exception of the backing PostgreSQL database) now use a single port (11371 by default).

Known Issues and Workarounds in Cloudera Navigator Key Trustee Server



Warning:

Interrupting deposit migration from Key Trustee Server to Key HSM can result in lost data

Workaround: Do not interrupt deposit migration to Key HSM.

Key Trustee Server active database setup fails when command line configuration specifies non-default database port

If the Key Trustee Server is configured from the command line to use a non-default database port (the default port is 11381), then when the Key Trustee Server service is added to Cloudera Manager, the first database startup fails.

Affected Versions: 5.15.0, 5.16.0

Cloudera Bug: KT-6238

Workaround:

1. Log in to the Key Trustee Server and manually stop/start the databases:

```
service keytrustee-db stop
```

Make sure that the `postmaster.pid` file, which is located in the database directory, no longer exists. If necessary, replace `/var/lib/keytrustee/db` in the following command with the appropriate database directory for your system:

```
# ls /var/lib/keytrustee/db/postmaster.pid
```

If `postmaster.pid` has not been cleaned up, then use the `pg_ctl` utility to stop the database directly:

```
# pg_ctl stop
```

2. Return to the Cloudera Manager home page, where the Key Trustee Server service is listed.
3. Go into the Key Trustee Server service configuration and for **Key Trustee Server Database Port**, specify the port that was configured during the command line configuration.
4. Redeploy the Key Trustee Server configuration and restart the service.

Cloudera Navigator Key Trustee Server Release Notes

Key Trustee Server 5.4.x and lower is not supported by Cloudera Manager 5.10.x and higher

Key Trustee Server 5.4.x and lower is not supported by Cloudera Manager 5.10.x and higher.

Affected Versions: 5.4.x

Fixed in Version(s): 5.10.1 through 5.14.0

Cloudera Bug: KT-6178

Workaround: Upgrade to a version of Key Trustee Server higher than 5.4.x.

Malicious user can access database used to store HDFS encryption zone keys

This vulnerability, which impacts Key Trustee Server, allows anyone with access to the host on which Key Trustee Server is running to use a non-libpq based connection driver to connect to the Postgres database used by Key Trustee Server without specifying a password.

The keys stored in the database are encrypted, which means access to data cannot be gained by this method. However, the vulnerability could provide a way for someone to alter encryption zone key metadata, delete the encryption zone keys, or delete the contents of the entire database.

Affected Versions: 5.3.0, 5.4.0, 5.5.0, 5.6.0, 5.7.0, 5.8.0, 5.9.0, 5.10.0, 5.11.0, 5.12.0, 5.13.0

Fixed in Version: 5.14.0

Cloudera Bug: DOCS-2788

Workaround: To remove the connection capability to Postgres (via the local network from all Key Trustee Server roles):

1. Log in as root to the host running Key Trustee Server.
2. Make a backup of the original Postgres configuration file:

```
cp /var/lib/keytrustee/db/pg_hba.conf \ /var/lib/keytrustee/db/pg_hba.conf.bak
```

3. Edit the Postgres configuration file to comment out the use of network connections for Postgres.
4. Using a preferred editor, edit `/var/lib/keytrustee/db/pg_hba.conf`, and add a # character to the beginning of the following lines in the file:

```
host keytrustee keytrustee 127.0.0.1/32 md5
host keytrustee keytrustee ::1/128 md5
```

After editing, the files should appear as:

```
#host keytrustee keytrustee 127.0.0.1/32 md5
#host keytrustee keytrustee ::1/128 md5
```

5. Save the file.
6. Perform a rolling restart of Key Trustee Server service.



Note:

While Key Trustee Server is restarting, new keys cannot be created and keys cannot be rolled. The restart should complete in under five minutes.

If there is only one instance of Key Trustee Server, restarting the Key Trustee Service may result in delayed or failed jobs waiting on an encryption zone key. If a job fails, rerun the job after the single instance of Key Trustee Server has restarted.

The encryption wizard continues to fail if there is a failure during initial configuration run

The encryption wizard continues to fail if there was a failure during the initial run configuring HSM KMS.

Affected Versions: 5.12.0, 5.13.0, 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-4909

Workaround: Open Cloudera Manager in another browser tab, and manually stop the installed KMS by clicking the arrow next to the KMS and selecting **Stop**. Then retry the installation in the new tab after correcting the cause of the install failure.

Before installing the Thales backed HSM KMS, you must add the KMS user to the `nfast` group

After installation of the Thales HSM client, and before installing Navigator HSM KMS backed by Thales HSM, you must add the KMS user to the `nfast` group..

Affected Versions: 5.12.0, 5.13.0, 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-4618

Workaround: Run the following command to manually add the KMS user to the `nfast` group:
`usermod -a -G nfast kms`

Navigator Key HSM may become unresponsive when there are network delays or connection issues with the HSM

This issue manifests itself as job failures within CDH. To verify that this issue is the source of the failed CDH jobs and that Key HSM has become unresponsive log into the Key HSM server and look at the log file in `/var/log/keyhsm/`. If no new entries are being written to the log file while key requests are being made, e.g. CDH jobs continue to retry, then follow the directions in the Workaround section.

Workaround: If threads accumulate and Navigator Key HSM has become unresponsive then restart Key HSM by executing the following commands:

```
$ sudo service keyhsm shutdown
$ sudo service keyhsm start
```

Key Trustee KMS cannot connect to Key Trustee Server using TLS versions other than 1.0 on JDK 7

If you have configured Key Trustee Server to use a TLS version other than 1.0, Key Trustee KMS fails to connect to Key Trustee Server, and key operations fail when using JDK 7.

Workaround: Use TLS version 1.0 only, or JDK 8.

Key Trustee Server cannot use TLS version 1.2 on RHEL 6

Configuring Key Trustee Server to use TLS version 1.2 causes Key Trustee Server to be unable to start.

Workaround: Use your operating system package manager to upgrade the `pyOpenSSL` package to version 1.4 or higher, or do not configure Key Trustee Server to use TLS version 1.2.

Upgraded passive Key Trustee Server fails to start due to incorrect ownership of `recovery.conf`

Passive Key Trustee Servers upgraded from Key Trustee Server 3.8.x or lower fail to start with the following error:

```
WARNING:root:stdout pg_basebackup: directory "/var/lib/pgsql/9.3/keytrustee" exists but
is not empty
:
Traceback (most recent call last):
  File "/usr/bin/ktadmin", line 484, in <module>
    main()
  File "/usr/bin/ktadmin", line 473, in main
    init_slave()
  File "/usr/bin/ktadmin", line 349, in init_slave
    pgsetup.base_backup(ARGS.pg_rootdir, ARGS.master, PKG, run_as=ARGS.postgres_user)
  File "/usr/lib/python2.6/site-packages/keytrustee/server/setup/postgres.py", line 206,
in base_backup
    run([pg_basebackup, '-D', dest, '--host=%s' % master_ip, '--port=%d' % port,
```

Cloudera Navigator Key Trustee Server Release Notes

```
'--username=%s' % db_user], run_as=run_as)
File "/usr/lib/python2.6/site-packages/keytrustee/util.py", line 145, in run
raise subprocess.CalledProcessError(p.returncode, cmd)
subprocess.CalledProcessError: Command '['/usr/pgsql-9.3/bin/pg_basebackup', '-D',
'/var/lib/pgsql/9.3/keytrustee', '--host=kt01.example.com', '--port=5432',
'--username=keytrustee']' returned non-zero exit status 1
```

Workaround: Change the owner and group of `/var/lib/pgsql/9.3/keytrustee/recovery.conf` to postgres:

```
$ sudo chown postgres:postgres /var/lib/pgsql/9.3/keytrustee/recovery.conf
```

Key Trustee Server PKCS8 private key cannot communicate with Key HSM

If its private key is in PKCS8 format, Key Trustee Server cannot communicate with Key HSM.

Cloudera Bug: KT-3172

Workaround: Convert the Key Trustee Server private key to raw RSA format.

Key Trustee Server backup script fails if PostgreSQL versions lower than 9.3 are installed

If PostgreSQL versions lower than 9.3 are installed on the Key Trustee Server host, the `ktbackup.sh` script fails with an error similar to the following:

```
pg_dump: server version: 9.3.11; pg_dump version: 9.2.14
pg_dump: aborting because of server version mismatch
```

Workaround: Uninstall the lower PostgreSQL version.

Issues Fixed in Cloudera Navigator Key Trustee Server

The following sections describe issues fixed in each Cloudera Navigator Key Trustee Server release.



Note: Releases 5.4.3 and 5.4.9 are the only maintenance releases of Key Trustee Server 5.4.

Release 5.5.2 is the only maintenance release of Key Trustee Server 5.5.

There is no Release 5.6 of Key Trustee Server.

There are no maintenance releases of Key Trustee Server 5.7.

There are no maintenance releases of Key Trustee Server 5.8.

There are no maintenance releases of Key Trustee Server 5.9.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.15.0

`ktbackup.sh` script not rolling GPG encrypted files

Encrypted backup tarballs have an additional `.gpg` file extension after `.tar` to indicate that they are encrypted (`.tar.gpg`). When the `ktbackup.sh` script searched for files to roll, it only sought files that ended in `.tar`, so it was not rolling the encrypted backup tarballs. The script has been updated to account for the `.tar.gpg` extension for encrypted backups, and now rolls them as expected.

Cloudera Bug: KT-6107

Issues Fixed in Cloudera Navigator Key Trustee Server 5.14.0

Malicious user can access database used to store HDFS encryption zone keys

This vulnerability, which impacted Key Trustee Server, allowed anyone with access to the host on which Key Trustee Server is running to use a non-libpq based connection driver to connect to the Postgres database used by Key Trustee Server without specifying a password.

The keys stored in the database are encrypted, which means access to data could not be gained by this method. However, the vulnerability could provide a way for someone to alter encryption zone key metadata, delete the encryption zone keys, or delete the contents of the entire database.

Cloudera Bug: DOCS-2788

Issues Fixed in Cloudera Navigator Key Trustee Server 5.13.0

There are no fixed issues in Cloudera Navigator Key Trustee Server 5.13.0.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.12.0

There are no fixed issues in Cloudera Navigator Key Trustee Server 5.12.0.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.11.1

There are no fixed issues in Cloudera Navigator Key Trustee Server 5.11.1.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.11.0

Key Trustee Server logs omit timestamps

The Key Trustee Server log entries do not include timestamps.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.10.1

Key Trustee Server logs omit timestamps

The Key Trustee Server log entries do not include timestamps.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.10.0

For new features in Key Trustee Server 5.10.0, see [New Features and Changes in Cloudera Navigator Key Trustee Server](#) on page 2785.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.9.0

For new features in Key Trustee Server 5.9.0, see [New Features and Changes in Cloudera Navigator Key Trustee Server](#) on page 2785.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.8.0

Key Trustee Server logs superfluous error during registration

When you register a client to the Key Trustee Server, the following error message appears in the Key Trustee Server log:

```
Traceback (most recent call last):
  File "/usr/lib/python2.6/site-packages/keytrustee/server/webapp.py", line 2151, in
lookup
    return hkp.lookup(g.hkpdb, app.config['LOCAL_FINGERPRINT'])
  File "/usr/lib/python2.6/site-packages/keytrustee/server/hkp.py", line 72, in lookup
```

Cloudera Navigator Key Trustee Server Release Notes

```
raise wzex.BadRequest("Operation not specified.")
BadRequest: 400: Operation not specified.
```

Issues Fixed in Cloudera Navigator Key Trustee Server 5.7.0

Email messages from Key Trustee Server use hostname instead of fully qualified domain name

Email sent from Key Trustee Server includes a link that uses the Key Trustee Server short name instead of the fully qualified domain name.

Host Inspector displays the wrong version of Key Trustee Server

Host Inspector displays the wrong Key Trustee Server version, showing the Key Trustee KMS version instead.

Key Trustee Server supports weak RC4 ciphers for TLS

Key Trustee Server supports weak RC4 ciphers for TLS connections, potentially allowing clients to negotiate a weaker connection.

The `ktadmin keyhsm` command fails with M2Crypto error

Running the `ktadmin keyhsm --server https://khs01.example.com:9090 --trust` command fails with the following error:

```
M2Crypto.X509.X509Error: 140405990356800:error:0906D06C:PEM routines:PEM_read_bio:no
start line:pem_lib.c:703:Expecting: CERTIFICATE
```

Issues Fixed in Cloudera Navigator Key Trustee Server 5.5.2

Key Trustee Server with password-protected private key cannot communicate with Key HSM

If its private key is password-protected, Key Trustee Server cannot communicate with Key HSM.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.5.0

Key Trustee Server missing from Components page

Key Trustee Server is not listed in the **Components** page (**Hosts > *Hostname* > Components**) of the host it is installed on.

Passive Database role fails to restart after enabling synchronous replication

After enabling synchronous replication, the Passive Database role fails to start when restarting the Key Trustee Server service in Cloudera Manager.

Key Trustee Server using CherryPy allows SSLv3

Key Trustee Server using CherryPy as the web backend allows SSLv3 connections, which are considered insecure.

Key Trustee Server sends wrong command in email notification

When you register a client, Key Trustee Server sends an email notification with a command to import the Key Trustee Server GPG key. The command uses the `hkps` protocol and port 80. The correct protocol is `hkps` and the correct port is 11371.

Enabling synchronous replication fails if `python-ordereddict` is not installed

On Key Trustee Server hosts where the `python-ordereddict` package is not installed, enabling synchronous replication fails, and the Passive Database role fails to start. The `stderr.log` file for the `DB_ACTIVE-setup` process (`/var/run/cloudera-scm-agent/process/ID-keytrustee_server-DB_ACTIVE-setup/logs/stderr.log`) contains the following error:

```
Traceback (most recent call last):
  File "aux/prop2x.py", line 14, in <module>
    from ordereddict import OrderedDict
ImportError: No module named ordereddict
```

Key Trustee Server does not respond to requests

On systems with low entropy, local fingerprint generation fails during setup, but Cloudera Manager reports no issues with the Key Trustee Server service. Subsequent attempts to use Key Trustee Server fail with an error similar to the following in the Key Trustee Server log:

```
2015-06-03 11:44:17,232 - keytrustee.server.webapp - ERROR - 'None' is not a valid
qualified fingerprint
Traceback (most recent call last):
  File
"/opt/cloudera/parcels/KEYTRUSTEE_SERVER-5.4.0-1.keytrustee5.4.0.p0.204/lib/python2.6/
site-packages/keytrustee-5.4.0-py2.6.egg/keytrustee/server/webapp.py", line 1917, in
index
    server.response = server.respond()
  File
"/opt/cloudera/parcels/KEYTRUSTEE_SERVER-5.4.0-1.keytrustee5.4.0.p0.204/lib/python2.6/
site-packages/keytrustee-5.4.0-py2.6.egg/keytrustee/server/webapp.py", line 1783, in
respond
    return Response(self.get_fingerprint(), **g.resp_kwargs)
  File
"/opt/cloudera/parcels/KEYTRUSTEE_SERVER-5.4.0-1.keytrustee5.4.0.p0.204/lib/python2.6/
site-packages/keytrustee-5.4.0-py2.6.egg/keytrustee/server/webapp.py", line 1338, in
get_fingerprint
    key_info = g.gpg.get_key_info(str(app.config['LOCAL_FINGERPRINT']))
  File
"/opt/cloudera/parcels/KEYTRUSTEE_SERVER-5.4.0-1.keytrustee5.4.0.p0.204/lib/python2.6/
site-packages/keytrustee-5.4.0-py2.6.egg/keytrustee/gnupg.py", line 415, in get_key_info

    search_string = fp_or_id(fingerprint, keyid)
  File
"/opt/cloudera/parcels/KEYTRUSTEE_SERVER-5.4.0-1.keytrustee5.4.0.p0.204/lib/python2.6/
site-packages/keytrustee-5.4.0-py2.6.egg/keytrustee/gnupg.py", line 817, in fp_or_id
    result = unqualified_fp(fingerprint)
  File
"/opt/cloudera/parcels/KEYTRUSTEE_SERVER-5.4.0-1.keytrustee5.4.0.p0.204/lib/python2.6/
site-packages/keytrustee-5.4.0-py2.6.egg/keytrustee/gnupg.py", line 761, in unqualified_fp

    _, _, fp = read_qualified_fp(fp)
  File
"/opt/cloudera/parcels/KEYTRUSTEE_SERVER-5.4.0-1.keytrustee5.4.0.p0.204/lib/python2.6/
site-packages/keytrustee-5.4.0-py2.6.egg/keytrustee/gnupg.py", line 770, in
read_qualified_fp
    raise ValueError("%s' is not a valid qualified fingerprint" % (s))
ValueError: 'None' is not a valid qualified fingerprint
```

Cannot register a client with Key Trustee Server after registering GPG public key

After creating an organization, the organization administrator receives a link to register a GPG key. If the administrator registers the GPG key before any clients are registered with the organization, all subsequent registrations with that organization fail.

Cloudera Navigator Key Trustee Server Release Notes

Starting Key Trustee Server with `--daemonize` option fails

Starting Key Trustee Server with the command `keytrustee-server start --daemonize` fails with the following error:

```
2015-04-01 05:33:26,843 - cherryypy.error - ERROR - [01/Apr/2015:05:33:26] ENGINE Error
in 'start' listener <bound method Daemonizer.start of
<cherryypy.process.plugins.Daemonizer object at 0x1e71490>>
Traceback (most recent call last):
  File "/usr/lib/python2.6/site-packages/cherryypy/process/wspbus.py", line 197, in
publish
    output.append(listener(*args, **kwargs))
  File "/usr/lib/python2.6/site-packages/cherryypy/process/plugins.py", line 380, in
start
    si = open(self.stdin, "r")
IOError: [Errno 2] No such file or directory: '/tmp/stdout'
```

Starting Key Trustee Server when it is already running updates `keytrustee.pid` file

If you try to start Key Trustee Server when it is already running, the `/var/lib/keytrustee/.keytrustee/keytrustee.pid` file is updated with the new process ID (PID). The new process fails to start, but the `keytrustee.pid` retains the new PID, and `service` commands fail with the following error:

```
Traceback (most recent call last):
  File "/usr/bin/keytrustee-server", line 136, in <module>
    main()
  File "/usr/bin/keytrustee-server", line 131, in main
    status()
  File "/usr/bin/keytrustee-server", line 104, in status
    pid_file = ProcessIDInfo(ARGS.pidfile)
  File "/usr/lib/python2.6/site-packages/keytrustee/process.py", line 8, in __init__
    raise ex.KeytrusteeError('Invalid PID path: %s' % path)
keytrustee.exceptions.KeytrusteeError: Invalid PID path:
/var/lib/keytrustee/.keytrustee/keytrustee.pid
```

Starting Key Trustee Server while the port is in use fails silently

If you try to start Key Trustee Server while the port is in use (for example, due to a stale process or another process using the port), the process appears to start successfully, but ends shortly after with no indication to the user. The Key Trustee Server log contains the following error:

```
Traceback (most recent call last):
  File "/usr/lib/python2.6/site-packages/cherryypy/process/wspbus.py", line 235, in start
    self.publish('start')
  File "/usr/lib/python2.6/site-packages/cherryypy/process/wspbus.py", line 215, in
publish
    raise exc
ChannelFailures: IOError("Port 11371 not free on 'keytrustee01.example.com',")
```

Starting Key Trustee Server as a user other than `keytrustee` fails

Starting Key Trustee Server as a user other than `keytrustee` (for example, as `root` or another super user) fails.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.4.9

Cannot register a client with Key Trustee Server after registering GPG public key

After creating an organization, the organization administrator receives a link to register a GPG key. If the administrator registers the GPG key before any clients are registered with the organization, all subsequent registrations with that organization fail.

Cloudera Bug: KT-1894

Issues Fixed in Cloudera Navigator Key Trustee Server 5.4.3

Key Trustee Server logs error GnuPG operation exited with return code=2

The Key Trustee Server log (`keytrustee.log`) includes error messages such as the following:

```
2014-12-30 04:23:22,166 - keytrustee.gnupg - ERROR - GnuPG operation exited with return
code=2:
[GNUPG:] ENC_TO 9375BEFD0BEC8C12 1 0
[GNUPG:] GOOD_PASSPHRASE
[GNUPG:] BEGIN_DECRYPTION
[GNUPG:] PLAINTEXT 62 1419942201
[GNUPG:] PLAINTEXT_LENGTH 230
[GNUPG:] ERRSIG DF46C6427B4466F3 1 10 00 1419942201 9
[GNUPG:] NO_PUBKEY DF46C6427B4466F3
[GNUPG:] DECRYPTION_OKAY
[GNUPG:] END_DECRYPTION
```

Package conflict: `python-jinja2`

Installing or upgrading Key Trustee Server on RHEL fails on hosts where `python-jinja2` is installed.

The `keytrustee-orgtool add` command fails with timeout: timed out **error on CentOS 6.6**

Adding an organization with the `keytrustee-orgtool add` command on CentOS 6.6 fails with the following error:

```
Exception in thread Thread-1:
Traceback (most recent call last):
  File "/usr/lib64/python2.6/threading.py", line 532, in __bootstrap_inner
    self.run()
  File "/usr/lib64/python2.6/threading.py", line 484, in run
    self.__target(*self.__args, **self.__kwargs)
  File "/usr/lib/python2.6/site-packages/keytrustee/server/util.py", line 363, in
send_mail
    smtp = smtplib.SMTP('localhost', timeout=timeout)
  File "/usr/lib64/python2.6/smtplib.py", line 239, in __init__
    (code, msg) = self.connect(host, port)
  File "/usr/lib64/python2.6/smtplib.py", line 296, in connect
    (code, msg) = self.getreply()
  File "/usr/lib64/python2.6/smtplib.py", line 337, in getreply
    line = self.file.readline()
  File "/usr/lib64/python2.6/socket.py", line 450, in readline
    data = self._sock.recv(self._rbufsize)
timeout: timed out
```

Passive Database role sometimes fails to start when starting the Key Trustee Server service

Starting or restarting the Key Trustee Server service attempts to start the Active Database role and Passive Database role. If the Active Database has not completed starting up when the Passive Database attempts to start, the Passive Database fails to start.

For parcel-based installations, client and path environments must be manually configured

Users cannot use Key Trustee Server command-line utilities without configuring the path and setting environment variables.

Initializing high availability Key Trustee Servers fails if SSH communication fails

Initializing the active Key Trustee Server fails if the active Key Trustee Server cannot reach the passive Key Trustee Server over the default SSH port (22).

Cloudera Navigator Key Trustee Server Release Notes

Enabling synchronous replication for high availability Key Trustee Servers fails

Running the `ktadmin enable-synchronous-replication` command does not properly configure synchronous replication in Key Trustee Server 5.4.0.

Issues Fixed in Cloudera Navigator Key Trustee Server 5.4.0

Cannot configure Key HSM after upgrade

After upgrading Key Trustee Server and Key HSM, the `ktadmin keyhsm --server khsm01.example.com --trust` command fails with the following error:

```
TypeError: coercing to Unicode: need string or buffer, NoneType found
```


Cloudera Navigator Key HSM Release Notes

These release notes provide information on the new and changed features, known issues, and fixed issues for Cloudera Navigator Key HSM.

New Features and Changes in Cloudera Navigator Key HSM

The following sections describe what's new and changed in each Cloudera Navigator Key HSM release.

What's New in Cloudera Navigator Key HSM

The following sections describe what's new in each Cloudera Navigator Key HSM release.

What's New in Cloudera Navigator Key HSM 1.11.0

- Support for RHEL 7.4.
- Support for SafeNet KeySecure 8.7.0.

What's New in Cloudera Navigator Key HSM 1.10.0

- Support for RHEL 6.9.
- Support for CentOS 6.9.

What's New in Cloudera Navigator Key HSM 1.9.0

- There are no new features in Cloudera Navigator Key HSM 1.9.0.

What's New in Cloudera Navigator Key HSM 1.8.0

- Key HSM supports RHEL 6.8.
- The Key HSM keystore now uses a randomly-generated password. To specify a password, set `keyhsm.keystore.password.set` to `yes` in the `application.properties` file.

What's New in Cloudera Navigator Key HSM 1.7.0

- Key HSM can be started in debug mode with more verbose logging.
- Key HSM logging has been improved.
- A new HTTP endpoint has been added to test hardware security module (HSM) connectivity and functionality.

What's New in Cloudera Navigator Key HSM 1.6.0

- Log size and rollover limits are configurable.

A number of issues have also been fixed. See [Issues Fixed in Cloudera Navigator Key HSM 1.6.0](#) on page 2799.

What's New in Cloudera Navigator Key HSM 1.5.1

An issue has been fixed. See [Issues Fixed in Cloudera Navigator Key HSM 1.5.1](#) on page 2799.

What's New in Cloudera Navigator Key HSM 1.5.0

- Key HSM supports RHEL 7.
- Key HSM restricts key names to universally allowed character sets to prevent problems with key migration.

A number of issues have also been fixed. See [Issues Fixed in Cloudera Navigator Key HSM 1.5.0](#) on page 2799.

What's New in Cloudera Navigator Key HSM 1.4.0

- Key HSM can now be attached to a Key Trustee Server instance with existing deposits.

- Support for SafeNet KeySecure
- KeySecure can be configured for HTTPS during setup

A number of issues have also been fixed. See [Issues Fixed in Cloudera Navigator Key HSM 1.4.0](#) on page 2800.

Known Issues and Workarounds in Cloudera Navigator Key HSM

Key HSM Luna setup not showing the correct login status

When running the `keyhsm setup luna` command, you are prompted for the Luna HSM slot number and login password. Key HSM then attempts to log into the Luna HSM to verify these settings are correct. In some circumstances, the setup script reports that the login was successful, even if it failed.

Affected Version: 5.14.0

Fixed in Version: 5.16.0

Cloudera Bug: KT-6623

Workaround: Any incorrect settings will cause the Key HSM service to throw an exception upon startup and exit. If the Key HSM service does not start correctly, check the log for the message: "Unable to sign into Luna HSM server. Please rerun application with 'setup' option to configure password." If this message appears, re-run the `keyhsm setup luna` command, and enter the correct slot number and login password.

Roll key command throws an exception and cannot retrieve metadata for key

When using Key Trustee KMS with Key Trustee Server and Key HSM (backed by an HSM device), if there is significant (> 15 ms ping time) network latency between the Key Trustee Server and the HSM device, then EDEK generation errors can occur during the roll key operation. These errors manifest in the KMS log as errors on the `generateEncryptedKey` operation. The KMS will recover from these errors on its own, but they may represent a nuisance to the operator.

Affected Version: 5.14.0, 5.16.0

Cloudera Bug: KT-5646

Workaround: When these errors occur, you can use the `hadoop key list -metadata` command to confirm whether or not the key roll was successful, despite the error condition.

Keys with certain special characters cannot be migrated from Key Trustee Server to Key HSM

If any existing key names in Key Trustee Server use special characters other than hyphen (-), period (.), or underscore (_), or begin with non-alphanumeric characters, the migration to Key HSM fails.

Workaround: Decrypt any data using the affected key names, and re-encrypt it using a new key name without special characters, and retry the migration.

Upgrading Key HSM removes init script and binary

Upgrading Key HSM from 1.4.x to 1.5.x and higher removes the Key HSM init script and `/usr/bin/keyhsm` binary.

Workaround: Reinstall Key HSM:

```
$ sudo yum reinstall keytrustee-keyhsm
```

Key HSM cannot trust Key Trustee Server certificate if it has extended attributes

Key HSM cannot trust the Key Trustee Server certificate if it has extended attributes, and therefore cannot integrate with Key Trustee Server.

Workaround: Import the Key Trustee Server certificate to the Key HSM trust store using Java `keytool` instead of the `keyhsm trust` command.

Issues Fixed in Cloudera Navigator Key HSM

The following sections describe issues fixed in each Cloudera Navigator Key HSM release.

Issues Fixed in Cloudera Navigator Key HSM 1.11.0

Minor stability improvements.

Issues Fixed in Cloudera Navigator Key HSM 1.10.0

Minor stability improvements.

Issues Fixed in Cloudera Navigator Key HSM 1.9.0

Minor stability improvements.

Issues Fixed in Cloudera Navigator Key HSM 1.8.0

For new features in Key HSM 1.8.0, see [What's New in Cloudera Navigator Key HSM 1.8.0](#) on page 2797.

Issues Fixed in Cloudera Navigator Key HSM 1.7.0

Using Key HSM with SafeNet KeySecure over TLS fails

When Key HSM is integrated with SafeNet KeySecure over TLS, Key HSM stops communicating with KeySecure after a period of time.

Issues Fixed in Cloudera Navigator Key HSM 1.6.0

Key HSM fails to start if the HSM contains a large number of keys

If the HSM stores a large number (hundreds) of keys, Key HSM fails to start with a timeout error similar to the following:

```
SEVERE: Timeout attempting to start services.
All services available:                :[ Failed ]
```

Key HSM logs do not roll over

Key HSM logs do not roll over, resulting in a log file that perpetually grows larger.

Issues Fixed in Cloudera Navigator Key HSM 1.5.1

Upgrading Key HSM removes init script

Upgrading Key HSM from 1.4.x to 1.5.x removes the Key HSM init script and `/usr/bin/keyhsm` binary. With this fix, future upgrades (from 1.5.1 to a higher release) do not experience this issue. To resolve the issue after upgrading to 1.5.x, reinstall Key HSM (`yum reinstall keytrustee-keyhsm`).

Issues Fixed in Cloudera Navigator Key HSM 1.5.0

Interrupting key migration from Key Trustee Server to Key HSM can result in lost data

Interrupting key migration (for example, using `Ctrl+C`) when integrating Key Trustee Server with Key HSM can result in data loss.

Issues Fixed in Cloudera Navigator Key HSM 1.4.0

Inserting deposits to Key HSM with Luna HSM logs stack trace

Inserting deposits to Key HSM using a Luna HSM logs the following stack trace to `keyhsm.log`:

```
SEVERE: -----Extended reason-----
Nov 21, 2014 4:39:11 PM com.cloudera.app.display.logging.AppLogger logInfo
INFO: com.safenetinc.luna.LunaCryptokiException.ThrowNew(LunaCryptokiException.java:66)
com.safenetinc.luna.LunaAPI.CheckSessionState(Native Method)
com.safenetinc.luna.LunaSession.isLoggedIn(LunaSession.java:149)
com.safenetinc.luna.LunaSlot.isLoggedIn(LunaSlot.java:155)
com.safenetinc.luna.LunaSlotManager.isLoggedIn(LunaSlotManager.java:585)
[...]
org.apache.tomcat.util.net.NioEndpoint$SocketProcessor.doRun(NioEndpoint.java:1720)
org.apache.tomcat.util.net.NioEndpoint$SocketProcessor.run(NioEndpoint.java:1679)
java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145)
java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615)
org.apache.tomcat.util.threads.TaskThread$WrappingRunnable.run(TaskThread.java:61)
java.lang.Thread.run(Thread.java:745)
```

`service keyhsm trust` **only works with absolute path**

Running `service keyhsm trust cert_file` fails if `cert_file` is a relative path.

Key Trustee KMS Release Notes

These release notes provide information on the new and changed features, known issues, and fixed issues for Key Trustee KMS.

New Features in Key Trustee KMS

The following sections describe what's new in each Key Trustee KMS release.



Note: Release 5.4.3 is the only maintenance release of Key Trustee KMS 5.4.

Release 5.5.4 is the only maintenance release of Key Trustee KMS 5.5.

There is no Release 5.6 of Key Trustee KMS.

Releases 5.7.1 and 5.7.4 are the only maintenance releases of Key Trustee KMS 5.7.

Release 5.8.2 is the only maintenance release of Key Trustee KMS 5.8.

There is no maintenance release of Key Trustee KMS 5.9.

There is no maintenance release of Key Trustee KMS 5.10.

What's New in Key Trustee KMS

The following sections describe what's new in each Key Trustee KMS release.

What's New in Key Trustee KMS 5.16.0

There is a newly-documented feature that describes how to move a Key Trustee KMS proxy service role instance from an existing cluster host to another cluster host. For details, refer to [Migrating a Key Trustee KMS Server Role Instance to a New Host](#).

What's New in Key Trustee KMS 5.15.0

There is a new feature that provides logic to detect and warn users about a potential problem where the GPG private keys have not been properly synchronized across all Key Trustee KMS HA hosts. For details, refer to [Validating Private Key Synchronization \(Key Trustee KMS HA Only\)](#), [Enabling Key Trustee KMS High Availability](#), and [Key Trustee KMS Encryption Issues](#).

What's New in Key Trustee KMS 5.14.0

There is a new procedure for migrating from a Key Trustee KMS (KT KMS) to a Hardware Security Module KMS (HSM KMS). For details, refer to [Migrating from a Key Trustee KMS to an HSM KMS](#).

What's New in Key Trustee KMS 5.13.0

- Re-encryption of encrypted data encryption keys (EDEKs), which ensures the automated re-encryption of all datafiles, and helps companies avoid the cost in time and resource of copying, decrypting, and re-encrypting datafiles. New, or re-encrypted key versions apply to datafiles previously encrypted by earlier key versions.
- Support for RHEL 7.4
- Support for OEL 7.4

What's New in Key Trustee KMS 5.12.0

- The Key Trustee Parcel will no longer be released via archive.cloudera.com. The parcel will now be released via www.cloudera.com/downloads. Parcels already released on the archive site will continue to be available there.
- Support for RHEL 6.9

Key Trustee KMS Release Notes

- Support for OEL 6.9

What's New in Key Trustee KMS 5.11.0

- Support for SLES 12 SP2
- Support for Ubuntu 16.04

What's New in Key Trustee KMS 5.10.0

- Support for RHEL 7.3
- Support for OEL 7.3

What's New in Key Trustee KMS 5.9.0

- Support for RHEL 6.8
- Support for OEL 6.8
- Support for Debian 8.4
- Support for SLES 12 SP1

What's New in Key Trustee KMS 5.8.2

- An issue has been fixed. See [Issues Fixed in Key Trustee KMS 5.8.2](#) on page 2806.

What's New in Key Trustee KMS 5.8.0

- Support for Debian 8.2
- When adding the parcel-based Key Trustee KMS service for the first time, Cloudera Manager automatically backs up Key Trustee KMS locally.
- The Key Trustee KMS backup script (`ktbackup.sh`) adds a new option, `--roll`, which specifies the number of backups to retain.
- When Key Trustee Server is configured for high availability, adding the Key Trustee KMS service in Cloudera Manager automatically configures round robin DNS load balancing.

What's New in Key Trustee KMS 5.7.4

- An issue has been fixed. See [Issues Fixed in Key Trustee KMS 5.7.4](#) on page 2807.

What's New in Key Trustee KMS 5.7.1

- An issue has been fixed. See [Issues Fixed in Key Trustee KMS 5.7.1](#) on page 2807.

What's New in Key Trustee KMS 5.7.0

- Support for RHEL 7.2
- Support for OEL 5.11, 7.2
- Support for SLES 11 SP4
- Support for Debian 7.8
- A backup script (`ktbackup.sh`) is included with Key Trustee KMS.

What's New in Key Trustee KMS 5.5.4

- A number of issues have been fixed. See [Issues Fixed in Key Trustee KMS 5.5.4](#) on page 2807.

What's New in Key Trustee KMS 5.5.0

- Rolling restart works with Key Trustee KMS high availability.
- When running Key Trustee KMS in a highly available configuration, Cloudera Manager can automatically generate the load balancer URL.

What's New in Key Trustee KMS 5.4.3

- A passive Key Trustee Server can be added to the Key Trustee KMS configuration using Cloudera Manager.

- An issue has also been fixed. See [Issues Fixed in Key Trustee KMS 5.4.3](#) on page 2808.

Known Issues and Workarounds in Key Trustee KMS

CDH upgrade failure

When upgrading to Key Trustee KMS 5.15.0 or 5.15.1 from Key Trustee KMS 5.14.0 or lower, *and* performing a rolling restart (instead of a full restart), the first Key Trustee KMS instance to restart may fail to come up and present the error:

```
"Unable to verify private key match between KMS hosts. If the system has been recently upgraded, DO NOT TAKE FURTHER ACTION and contact your support representative as soon as possible. If this is a new installation, verify private key files have been synched between all KMS hosts. Aborting to prevent data inconsistency."
```

Affected Versions: 5.15.0, 5.15.1

Fixed in Version: 5.16.0

Cloudera Bug: KT-6499

Workaround: If possible, perform a full restart instead of a rolling restart.

If you cannot execute a full restart, then add the following line to the `/var/lib/kms-keytrustee/keytrustee/.keytrustee/keytrustee.conf` file on all Key Trustee KMS instances, and then restart the Key Trustee KMS that failed:

```
"FINGERPRINT_VALIDATED" : "True"
```

Error when enabling Key Trustee KMS HA mode

When a new Key Trustee KMS instance is added to an existing Key Trustee KMS service, the new instance may initially fail to start up and return the error message:

```
Unable to verify private key match between KMS hosts.
```

This error will be corrected automatically after the GPG private keys are synchronized and the KMS service is restarted to pick up the new configuration. No user intervention is required.

Affected Versions: 5.15.0

Fixed in Version: 5.16.0

Cloudera Bug: KT-6471

Workaround: None.

Validation fails if `hostname` command returns `shortname`

If the `hostname` command on the OS returns `shortname`, and the `core-site.xml` of the KMS process has a `hadoop.security.key.provider.path` with a fully qualified domain name (FQDN), then the `znodes` will be created with the `shortname`. Consequently, when KMS 1 checks the fingerprint of KMS 2, it will expect the FQDN as the `znode`, and fail the validation.

Affected Versions: 5.15.0

Fixed in Version: 5.15.1

Cloudera Bug: KT-6412

Workaround: In `/var/lib/kms-keytrustee/keytrustee/.keytrustee/keytrustee.conf`, change the value of `FINGERPRINT_VALIDATED` from `True` to `False`.

Incorrect default KMS ACL values allow remote access to purge and undelete API calls on encryption zone keys

The Navigator Key Trustee KMS includes two API calls in addition to those in Apache Hadoop KMS: `purge` and `undelete`. The KMS ACL values for these commands are `keytrustee.kms.acl.PURGE` and `keytrustee.kms.acl.UNDELETE`, respectively. The default value for the ACLs in Key Trustee KMS 5.12.0 and 5.13.0 is `"*`", which allows anyone who knows the name of an encryption zone key and has network access to the Key Trustee KMS to make those calls against known encryption zone keys. The `UNDELETE` command will result in the recovery of a previously deleted, but not purged, key. The key will still be protected by normal ACLs and is not exposed. The `PURGE` command will permanently delete a key in active use, resulting in loss of access to encrypted HDFS data.

Affected Versions:

- Navigator Key Trustee KMS 5.12.0, 5.13.0
- Cloudera Manager 5.12.0, 5.12.1, 5.12.2
- Cloudera Manager 5.13.0, 5.13.1

Fixed in Version: 5.14.0

Cloudera Bug: DOCS-2910

Workaround: Use Cloudera Manager to set the values of the ACLs for `purge` (`keytrustee.kms.acl.PURGE`) and `undelete` (`keytrustee.kms.acl.UNDELETE`) to be empty; this denies access to all calls to those functions.

To set the values:

1. Log into Cloudera Manager as an administrative user.
2. Click on the **Key Trustee KMS** service.
3. Click on **Configuration**.
4. In the Search box, type "ACL".
5. At the bottom of the section **Key Management Server Proxy Advanced Configuration Snippet (Safety Valve) for kms-acls.xml** add a new ACL by clicking the + icon.
6. In the **Name** field add "keytrustee.kms.acl.PURGE".
7. Repeat step 5, and in the **Name** field add "keytrustee.kms.acl.UNDELETE".
8. Scroll to the bottom of the page and click **Save**.
9. Use Cloudera Manager to perform a rolling restart of the Navigator Key Trustee KMS service.



Note: If there is only one instance of Navigator Key Trustee KMS, then restarting the service may result in delayed or failed jobs waiting on an encryption zone key. If a job fails, then rerun the job after the single instance of Navigator Key Trustee KMS has restarted.

Cannot re-encrypt an encryption zone if a previous re-encryption on it was canceled

When canceling a re-encryption on an encryption zone, the status of the re-encryption may continue to show "Processing". When this occurs, future re-encrypt commands for this encryption zone will fail inside the NameNode, and the re-encryption will never complete.

Affected Version: 5.13.0

Fixed Version: 5.13.1

Cloudera Bug: CDH-59073

Workaround: To halt, or remove the "Processing" status for the encryption zone, re-issue the cancel re-encryption command on the encryption zone. If a new re-encryption command is required for this encryption zone, restart the NameNode before issuing the command.

Adding Key Trustee KMS 5.4 to Cloudera Manager 5.5 displays warning

Adding the Key Trustee KMS service to a CDH 5.4 cluster managed by Cloudera Manager 5.5 displays the following message, even if Key Trustee KMS is installed:

"The following selected services cannot be used due to missing components: keytrustee-keyprovider. Are you sure you wish to continue with them?"

Affected Version: 5.4

Workaround: Verify that the Key Trustee KMS parcel or package is installed and click **OK** to continue adding the service.

The Key Trustee KMS service fails to start if the Trust Store is configured without also configuring the Keystore

If you configure the Key Trustee KMS service **Key Management Server Proxy TLS/SSL Certificate Trust Store File** and **Key Management Server Proxy TLS/SSL Certificate Trust Store Password** parameters without also configuring the **Key Management Server Proxy TLS/SSL Server JKS Keystore File Location** and **Key Management Server Proxy TLS/SSL Server JKS Keystore File Password** parameters, the Key Trustee KMS service does not start.

Workaround: Configure all Trust Store and Keystore parameters.

Key Trustee KMS backup script fails if PostgreSQL versions lower than 9.3 are installed

If PostgreSQL versions lower than 9.3 are installed on the Key Trustee KMS host, the `ktbackup.sh` script fails with an error similar to the following:

```
pg_dump: server version: 9.3.11; pg_dump version: 9.2.14
pg_dump: aborting because of server version mismatch
```

Workaround: Uninstall the lower PostgreSQL version.

Issues Fixed in Key Trustee KMS

The following sections describe issues fixed in each Key Trustee KMS release.



Note: Release 5.4.3 is the only maintenance release of Key Trustee KMS 5.4.

Release 5.5.4 is the only maintenance release of Key Trustee KMS 5.5.

There is no Release 5.6 of Key Trustee KMS.

Releases 5.7.1 and 5.7.4 are the only maintenance releases of Key Trustee KMS 5.7.

Release 5.8.2 is the only maintenance release of Key Trustee KMS 5.8.

There is no maintenance release of Key Trustee KMS 5.9.

There is no maintenance release of Key Trustee KMS 5.10.

Issues Fixed in Key Trustee KMS 5.16.0

CDH upgrade failure

When upgrading to Key Trustee KMS 5.15.0 or 5.15.1 from Key Trustee KMS 5.14.0 or lower, *and* performing a rolling restart (instead of a full restart), the first Key Trustee KMS instance to restart may fail to come up and present an error.

Cloudera Bug: KT-6499

Error when enabling Key Trustee KMS HA mode

When a new Key Trustee KMS instance is added to an existing Key Trustee KMS service, the new instance may initially fail to start up and return the error message:

```
Unable to verify private key match between KMS hosts.
```

Key Trustee KMS Release Notes

This error will be corrected automatically after the GPG private keys are synchronized and the KMS service is restarted to pick up the new configuration. No user intervention is required.

Cloudera Bug: KT-6471

Issues Fixed in Key Trustee KMS 5.15.1

Validation fails if `hostname` command returns `shortname`

If the `hostname` command on the OS returns `shortname`, and the `core-site.xml` of the KMS process has a `hadoop.security.key.provider.path` with a fully qualified domain name (FQDN), then the `znodes` will be created with the `shortname`. Consequently, when KMS 1 checks the fingerprint of KMS 2, it will expect the FQDN as the `znode`, and fail the validation.

Issues Fixed in Key Trustee KMS 5.15.0

For new features in Key Trustee KMS 5.15.0, see [What's New in Key Trustee KMS 5.15.0](#) on page 2801.

Issues Fixed in Key Trustee KMS 5.14.0

For new features in Key Trustee KMS 5.14.0, see [What's New in Key Trustee KMS 5.14.0](#) on page 2801.

Issues Fixed in Key Trustee KMS 5.13.1

Cannot re-encrypt an encryption zone if a previous re-encryption on it was canceled

When canceling a re-encryption on an encryption zone, the status of the re-encryption would remain at "Processing". When this occurred, future re-encrypt commands for this encryption zone failed inside the NameNode, and the re-encryption never completed. This issue has been fixed.

Issues Fixed in Key Trustee KMS 5.13.0

For new features in Key Trustee KMS 5.13.0, see [What's New in Key Trustee KMS 5.13.0](#) on page 2801.

Issues Fixed in Key Trustee KMS 5.12.0

For new features in Key Trustee KMS 5.12.0, see [What's New in Key Trustee KMS 5.12.0](#) on page 2801.

Issues Fixed in Key Trustee KMS 5.11.0

For new features in Key Trustee KMS 5.11.0, see [What's New in Key Trustee KMS 5.11.0](#) on page 2802.

Issues Fixed in Key Trustee KMS 5.10.0

For new features in Key Trustee KMS 5.10.0, see [What's New in Key Trustee KMS 5.10.0](#) on page 2802.

Issues Fixed in Key Trustee KMS 5.9.0

For new features in Key Trustee KMS 5.9.0, see [What's New in Key Trustee KMS 5.9.0](#) on page 2802.

Issues Fixed in Key Trustee KMS 5.8.2

KMS does not fail over to Passive Key Trustee Server in some network failure scenarios

In some situations, if the Active Key Trustee Server is unreachable on the network, Key Trustee KMS does not fail over to the Passive Key Trustee Server.

Cloudera Bug: KT-3289

Issues Fixed in Key Trustee KMS 5.8.0

For new features in Key Trustee KMS 5.8.0, see [What's New in Key Trustee KMS 5.8.0](#) on page 2802.

Issues Fixed in Key Trustee KMS 5.7.4

KMS does not fail over to Passive Key Trustee Server in some network failure scenarios

In some situations, if the Active Key Trustee Server is unreachable on the network, Key Trustee KMS does not fail over to the Passive Key Trustee Server.

Cloudera Bug: KT-3289

Issues Fixed in Key Trustee KMS 5.7.1

KMS ACLs read from wrong file

The UNDELETE and PURGE ACL entries were being read from `kms-site.xml` instead of `kms-acls.xml`.

Cloudera Bug: KT-2857

Issues Fixed in Key Trustee KMS 5.7.0

For new features in Key Trustee KMS 5.7.0, see [What's New in Key Trustee KMS 5.7.0](#) on page 2802.

Issues Fixed in Key Trustee KMS 5.5.4

Key Trustee KMS configuration file and keys are stored in a volatile location

If the Key Trustee KMS 5.5.0 parcel is deactivated, any existing GPG keys are also deactivated. If the parcel is then reactivated, new GPG keys (used to create an authenticated and private communication channel with the Key Trustee Server) are generated. The existing GPG keys that were in use before the deactivation are not lost; however, they become inactive. If remedial action is not taken before deactivation, this can result in a loss of access to HDFS Encryption Zone keys generated with the older set of GPG keys. This in turn leads to loss of access to all data in all encryption zones. As long as the Key Trustee KMS parcel directory is not deleted, access can be restored. Assistance from Cloudera Support may be required. See [TSB 2016-121](#) for more information (requires login to the Cloudera Support Portal).

Cloudera Bug: KT-2460

KMS ACLs read from wrong file

The UNDELETE and PURGE ACL entries were being read from `kms-site.xml` instead of `kms-acls.xml`.

Cloudera Bug: KT-2858

Issues Fixed in Key Trustee KMS 5.5.0

Key Trustee KMS intermittently fails to start due to permission issues

Key Trustee KMS sometimes fails to start with the following error:

```
java.io.IOException: TrusteeKeyProvider initialization failed.
    at
    com.cloudera.keytrustee.TrusteeKeyProvider.createInitialClient(TrusteeKeyProvider.java:212)
    at com.cloudera.keytrustee.TrusteeKeyProvider.<init>(TrusteeKeyProvider.java:114)
    at com.cloudera.keytrustee.TrusteeKeyProvider.<init>(TrusteeKeyProvider.java:86)
    [...]
Caused by: java.io.FileNotFoundException:
/var/lib/kms-keytrustee/keytrustee/.keytrustee/keytrustee.conf (Permission denied)
    at java.io.FileOutputStream.open(Native Method)
    at java.io.FileOutputStream.<init>(FileOutputStream.java:221)
```

Key Trustee KMS Release Notes

```
at java.io.FileOutputStream.<init>(FileOutputStream.java:171)
at
com.cloudera.keytrustee.impl.ClientFactoryImpl.createNewClient(ClientFactoryImpl.java:129)
... 30 more
```

Key Trustee KMS starts up successfully even if it cannot contact the Key Trustee Server

Key Trustee KMS starts successfully when it cannot connect to the Key Trustee Server, instead of properly failing to start.

Issues Fixed in Key Trustee KMS 5.4.3

[Host Component page display](#)

The Host Component page now displays the package version for the KMS Trustee Key Provider.

Cloudera Bug: OPSAPS-25692

Cloudera Navigator HSM KMS Release Notes

These release notes provide information on the new and changed features, known issues, and fixed issues for Cloudera Navigator HSM KMS.

New Features and Changes in Cloudera Navigator HSM KMS

The following sections describe what's new and changed in each Cloudera Navigator HSM KMS release.

What's New in Cloudera Navigator HSM KMS

The following section describes what's new in this Cloudera Navigator HSM KMS release.

What's New in Cloudera Navigator HSM KMS 5.15.0

This release of Cloudera Navigator HSM KMS does not include new features or functionality.

What's New in Cloudera Navigator HSM KMS 5.14.0

This release of Cloudera Navigator HSM KMS provides the following new functionality:

- When an enterprise requires the highest level of encryption zone key (EZ key) security, or, it must adhere to strict security certification compliance standards, users now have the option of migrating from a Key Trustee KMS (KT KMS) to a Hardware Security Module KMS (HSM KMS). Migrating from KT KMS to HSM KMS essentially means that EZ key metadata is moved from one KMS to another. For details, see [Migrating from a Key Trustee KMS to an HSM KMS](#).

What's New in Cloudera Navigator HSM KMS 5.13.0

This release of Cloudera Navigator HSM KMS provides the following new functionality:

- Support for RHEL 7.4.

What's New in Cloudera Navigator HSM KMS 5.12.0

This is the first release of the new Cloudera Navigator HSM KMS. The new component provides:

- HSM KMS backed by Thales HSM is a custom Key Management Server (KMS) that uses a supported Thales HSM as the underlying keystore, instead of the file-based Java KeyStore (JKS) used by the default Hadoop KMS.
- Navigator HSM KMS backed by Luna HSM is a custom Key Management Server (KMS) that uses a supported Luna HSM as the underlying keystore, instead of the file-based Java KeyStore (JKS) used by the default Hadoop KMS.

Known Issues and Workarounds in Cloudera Navigator HSM KMS

HSM KMS `get_keys` returns wrong number of keys even though creation of one key failed

If the HSM KMS is left running while its local HSM KMS Metastore is stopped, then commands that attempt to write to the metastore (such as `create key`, `roll key`, and `delete key`) will fail (expected behavior). However, commands that retrieve key metadata and/or key material (such as `key list`) may appear to succeed, but can return incorrect results (for example, an empty key list when there are extant keys) due to cache invalidation without the ability to refresh the cache from the metastore.

Affected Versions: 5.12.0, 5.13.0, 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-5060

Workaround: None

"Exception in generateEncryptedKeys" error in HSM KMS log file after key migration

During the migration of encryption zone keys from a KT KMS to an HSM KMS, you may see errors related to the `generateEncryptedKeys` call (this operation populates the EDEK cache) on the HSM KMS node that is not running the migration. These errors will not prevent the migration operation from completing and will not cause ongoing problems after migration completes.

Affected Version: 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-5672

Workaround: None

KT KMS migration to HSM KMS fails if a key is not found on HSM

When using Key Trustee KMS with Key Trustee Server and Key HSM (backed by an HSM device), if a key version is deleted on the HSM device directly without also deleting and purging that key version on the Key Trustee KMS, then attempts to migrate from the KT KMS to the HSM KMS will fail.

Affected Version: 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-5671

Workaround: None

Timeout error during encryption zone key creation

There are situations where the key cache is synchronously populated to capacity during the create encryption zone operation. The expected behavior is that the key cache is synchronously populated only to the low watermark level (the rest of the keys should be created asynchronously).

Affected Version: 5.13.0, 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-5296

Workaround: On the HSM KMS, in the field **HSM KMS Proxy Advanced Configuration Snippet (Safety Valve)** for `kms-site.xml`:

- `hadoop.security.kms.encrypted.key.cache.low.watermark .05`
- `hadoop.security.kms.encrypted.key.cache.size 30`

On the HDFS, in the field **HDFS Cluster-wide Advanced Configuration Snippet (Safety Valve)** for `core-site.xml`:

- `hadoop.security.kms.client.encrypted.key.cache.size 30`
- `hadoop.security.kms.client.encrypted.key.cache.low-watermark .05`

Key description is not synchronized with the second metastore

If the `description` option is specified as the argument in a Hadoop `key create` command, then the description information is stored only on the KMS instance that responds to the create request. The description metadata is not synchronized to the other KMS instances in the role group. This does not affect normal key operations.

Affected Version(s): 5.12.0

Fixed Version: 5.12.1

Cloudera Bug: KT-5042

Workaround: Use the `-provider` argument on the key list operation to target key queries to a specific KMS instance.

HSM KMS Luna may need to be restarted if inactive for extended period

If Hadoop key operations return `com.safenetinc.luna.exception.LunaCryptokiException` after the KMS has been running without activity for an extended period time, the Luna session may have been dropped.

Affected Version(s): 5.12.0, 5.13.0, 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-5018

Workaround: Restart the KMS service.

[Creating multiple instances of HSM KMS on the same host and port causes an error upon delete](#)

Creating a KMS role instance on a host that previously hosted a KMS role instance in the same role group that had its data directories deleted results in errors when attempting to run Hadoop `key delete` operations.

Affected Version(s): 5.12.0, 5.13.0, 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-4992

Workaround: This workaround requires the assistance of Cloudera support; request assistance with issue KT-4992

[Incorrect status for "Restart stale services" step in HDFS encryption wizard post-service installation](#)

There are times when completion of the HDFS Encryption Wizard does not show an active "Restart stale services and redeploy client configuration" link.

Affected Version(s): 5.12.0, 5.13.0, 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-4987

Workaround: Refresh the page and the link should become active.

[The encryption wizard continues to fail if there is a failure during initial configuration run](#)

The encryption wizard continues to fail if there was a failure during the initial run configuring HSM KMS.

Affected Version(s): 5.12.0, 5.13.0, 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-4909

Workaround: Open Cloudera Manager in another browser tab, and manually stop the installed KMS by clicking the arrow next to the KMS and selecting **Stop**. Then retry the installation in the new tab after correcting the cause of the install failure.

[Before installing the Thales backed HSM KMS, you must add the KMS user to the `nfast` group](#)

After installation of the Thales HSM client, and before installing Navigator HSM KMS backed by Thales HSM, you must add the KMS user to the `nfast` group..

Affected Version(s): 5.12.0, 5.13.0, 5.14.0, 5.15.0, 5.16.0

Cloudera Bug: KT-4618

Workaround: Run the following command to manually add the KMS user to the `nfast` group:
`usermod -a -G nfast kms`

Issues Fixed in Cloudera Navigator HSM KMS

The following section describe issues fixed in each HSM KMS release.

Issues Fixed in Cloudera Navigator HSM KMS 5.15.0

For details about what is new in this release, see [What's New in Cloudera Navigator HSM KMS 5.15.0](#) on page 2809.

Issues Fixed in Cloudera Navigator HSM KMS 5.14.0

For details about what is new in this release, see [New Features and Changes in Cloudera Navigator HSM KMS](#) on page 2809.

Issues Fixed in Cloudera Navigator HSM KMS 5.13.0

Failure creating keys after Metastore restart.

Previously, stopping or restarting the HSM KMS Metastore role while the HSM KMS Proxy role remained running resulted in stale metastore connections and errors such the following: "FATAL: terminating connection due to administrator command."

Cloudera Bug: KT-4912

Issues Fixed in Cloudera HSM KMS 5.12.1

Key version number information now reported correctly by Hadoop `key list -metadata` operation after key roll.

Previously, after the key roll operation, the key's version number remained the same. Now, rather than returning a hard-coded version number on keys loaded from metastore, that information is returned dynamically.

Cloudera Bug: KT-5044

Key description is not synchronized with the second metastore

When the description option was specified as the argument in a Hadoop `key create` command, then the description information was stored only on the KMS instance that responded to the create request. The description metadata was not synchronized to the other KMS instances in the role group. This issue did not affect normal key operations.

Cloudera Bug: KT-5042

Issues Fixed in Cloudera Navigator HSM KMS 5.12.0

This is the first release of Cloudera Navigator HSM KMS and consequently, there are no fixed issues. For details about what is new in this release, see [New Features and Changes in Cloudera Navigator HSM KMS](#) on page 2809.

Cloudera Navigator Encrypt Release Notes

These release notes provide information on the new and changed features, known issues, and fixed issues for Cloudera Navigator Encrypt.

New Features and Changes in Cloudera Navigator Encrypt

The following sections describe what's new and changed in each Cloudera Navigator Encrypt release.

What's New in Cloudera Navigator Encrypt

The following sections describe what's new in each Cloudera Navigator Encrypt release.

What's New in Cloudera Navigator Encrypt 3.16.0

- Support for RHEL/CentOS 7.5 including kernel version 3.10.0-862.14.4

What's New in Cloudera Navigator Encrypt 3.15.1

- Support for RHEL/CentOS 7.5 up to but not including kernel version 3.10.0-862.14.4

What's New in Cloudera Navigator Encrypt 3.15.0

- Support for SLES 12 SP3

What's New in Cloudera Navigator Encrypt 3.14.0

- The `navencrypt exec` command option has been deprecated.
- Support for RHEL 7.4
- Support for OEL 7.4

What's New in Cloudera Navigator Encrypt 3.13.0

- Support for RHEL 6.9
- Support for CentOS 6.9
- Navigator Encrypt now supports preparing devices using a UUID, in addition to device name. This UUID is simply a symbolic link to the actual device, and is created when preparing a device with Navigator Encrypt during a `navencrypt-prepare` operation.

What's New in Cloudera Navigator Encrypt 3.12.0

- Support for SLES 12 SP2
- Support for Ubuntu 16.04

What's New in Cloudera Navigator Encrypt 3.11.0

- Support for RHEL 6.8, 7.3
- Support for SLES 12 SP1

What's New in Cloudera Navigator Encrypt 3.10.0

- Support for OEL 7.2
- Several new commands have been added:
 - `navencrypt-collect` displays environment information for troubleshooting. For more information, see [Collecting Navigator Encrypt Environment Information](#).
 - `navencrypt-move --list-categories` lists existing categories.

Cloudera Navigator Encrypt Release Notes

- `navencrypt restore-control-file` restores the `/etc/navencrypt/control` file from Key Trustee Server.
- `navencrypt status --integrity` validates the Navigator Encrypt configuration.
- Access control list (ACL) policy files now support comments.

What's New in Cloudera Navigator Encrypt 3.9.0

- Support for RHEL 7.2
- Support for SLES 11 SP4
- Support for Debian 7.8
- Support for Ubuntu 14.04.3

An issue has also been fixed. See [Issues Fixed in Cloudera Navigator Encrypt 3.9.0](#) on page 2817.

What's New in Cloudera Navigator Encrypt 3.8.0

- Support for RHEL 7.1
- Support for automatically mounting loop devices on boot
- Automatic failover for high availability Key Trustee Servers
- The ability to restore deleted local mount encryption keys (MEKs) from Key Trustee Server by UUID

What's New in Cloudera Navigator Encrypt 3.7.1

- A number of issues have been fixed. See [Issues Fixed in Cloudera Navigator Encrypt 3.7.1](#) on page 2818.

What's New in Cloudera Navigator Encrypt 3.7.0

- Migration utility to migrate zNcrypt to Cloudera Navigator Encrypt

What's Changed in Cloudera Navigator Encrypt

The following sections describe what's changed in each Cloudera Navigator Encrypt release.

What's Changed in Cloudera Navigator Encrypt 3.8.0

- The kernel module is renamed from `zncrypt-kernel-module` to `navencrypt-kernel-module`.

Known Issues and Workarounds in Cloudera Navigator Encrypt

Navigator Encrypt will not build on RHEL kernel 3.10.0-862.14.4

The Navigator Encrypt kernel module will not build on RHEL kernel 3.10.0-862.14.4. This impacts new installations, and existing installations that are upgrading to kernel 3.10.0-862.14.4. This issue prevents the `navencrypt-mount` service from running and Navigator Encrypt mount points from being accessible.

Affected Versions: 3.15.1

Fixed in Version: 3.16.0

Cloudera Bug: KT-6677

Workaround: Upgrade to Navigator Encrypt 3.16.0 before upgrading to RHEL kernel 3.10.0-862.14.4.

Interactive console does not support UUID and datastore options

The Interactive console does not support UUID (3.13 and later) and datastore options (3.14.0 and earlier).

Affected Versions: Navigator Encrypt versions 3.6, 3.7, 3.8, 3.9, 3.10, 3.11, 3.12, 3.13, 3.14, 3.15, 3.16.

Fixed in Version: None

Cloudera Bug: DOCS-2505

Workaround: None.

The Navigator Encrypt `-exec` option does not work.

The `navencrypt -exec` command option does not work.

Affected Versions: Navigator Encrypt versions 3.6, 3.7, 3.8, 3.9, 3.10, 3.11, 3.12, 3.13

Fixed in Version: Navigator Encrypt 3.14.0 (option deprecated)

Cloudera Bug: KT-5113

Workaround: Execute a temporal command such as `cp /data/encrypted /destination/path` to temporarily add the process (for example, `cp`) to the ACLs.

Navigator Encrypt progress bar random display anomaly

When running the `navencrypt-move` command, the graphic display of progress randomly breaks. This display anomaly does *not* impact the move operation and there is no data loss. This is a graphic display issue only, and is not related to the `move` command operation being performed.

When the display issue occurs, it appears as follows:

```
Moving from: '/local/mnt/navencrypt/mount/mysql/local/mnt/mis/mysql/2net_kernel'
Moving to: '/local/mnt/mis/mysql/2net_kernel'
100% [====/usr/sbin/navencrypt-move: line 462: printf: write error: Interrupted system
call
100% [=====>] [ 72.06 GB]
```

Affected Versions: CDH: 5.3.0, 5.4.0, 5.5.0, 5.6.0, 5.7.0, 5.8.0, 5.9.0, 5.10.0, 5.11.0, 5.12.0/Navigator Encrypt: 3.6, 3.7, 3.8, 3.9, 3.10, 3.11, 3.12, 3.13

Fixed in Version: CDH 5.15.0/Navigator Encrypt 3.15

Cloudera Bug: KT-5093

Workaround: None.

The `navencrypt status --integrity` command does not detect loop devices properly

Running the `navencrypt status --integrity` command on a system using a loop device incorrectly reports that the backing file does not exist.

Affected Versions: CDH 5.8.0/Navigator Encrypt 3.10.

Fixed in Version: CDH 5.10.0/Navigator Encrypt 3.11.

Cloudera Bug: KT-3174

Workaround: None.

Navigator Encrypt 3.12 and earlier does not support specific kernels

Versions of Navigator Encrypt 3.12 and earlier do not support the following kernels:

- 3.10.0-514.2.2.el7.x86_64
- 3.10.0-514.16.1.el7.x86_64
- 3.10.0-514.10.2.el7.x86_64

Note that this is not specific to a particular platform; rather, it is specific only to the kernels listed here.

Affected Versions: CDH 5.9.0/Navigator Encrypt: 3.6, 3.7, 3.8, 3.9, 3.10, 3.11, 3.12

Fixed in Version: CDH 5.11.0/Navigator Encrypt 3.13

Cloudera Bug: KT-4727

Cloudera Navigator Encrypt Release Notes

Workaround: Use a kernel other than those listed here.

The Navigator Encrypt kernel module does not build on Ubuntu 14.04.04

The Navigator Encrypt kernel module does not build on Ubuntu 14.04.04 (kernel version 4.2).

Affected Versions: CDH 5.7.0/Navigator Encrypt 3.9

Fixed in Version: CDH 5.8.0/Navigator Encrypt 3.10

Cloudera Bug: KT-2544

Workaround: None.

Using `screen` utility with Cloudera Navigator Encrypt does not work

Running Cloudera Navigator Encrypt commands within the Linux `screen` utility does not work correctly.

Affected Versions: CDH: 5.3, 5.4.0

Fixed in Version: CDH 5.4.9, 5.5.0

Cloudera Bug: KT-1971

Workaround: None.

Sophos antivirus is not compatible with Cloudera Navigator Encrypt

Kernel panics have been observed on machines running both Sophos antivirus and Cloudera Navigator Encrypt.

Affected Versions: CDH: 5.3, 5.4, 5.5, 5.6

Fixed in Version: CDH 5.7.x

Cloudera Bug: KT-3098

Workaround: None. Do not use Sophos antivirus in conjunction with Cloudera Navigator Encrypt.

Cloudera Navigator Encrypt crashes in Docker container

Cloudera Navigator Encrypt crashes when run in a Docker container.

Affected Versions: CDH: 5.3, 5.4, 5.5, 5.6

Fixed in Version: CDH 5.7.x

Cloudera Bug: KT-3098

Workaround: None. Do not use Cloudera Navigator Encrypt with Docker.

Issues Fixed in Cloudera Navigator Encrypt

The following sections describe issues fixed in each Cloudera Navigator Encrypt release.

Issues Fixed in Cloudera Navigator Encrypt 3.16.0

See [New Features and Changes in Cloudera Navigator Encrypt](#) on page 2813 for new features in Cloudera Navigator Encrypt 3.16.0.

Issues Fixed in Cloudera Navigator Encrypt 3.15.0

See [New Features and Changes in Cloudera Navigator Encrypt](#) on page 2813 for new features in Cloudera Navigator Encrypt 3.15.0.

Issues Fixed in Cloudera Navigator Encrypt 3.14.0

See [New Features and Changes in Cloudera Navigator Encrypt](#) on page 2813 for new features in Cloudera Navigator Encrypt 3.14.0.

Issues Fixed in Cloudera Navigator Encrypt 3.13.0

See [New Features and Changes in Cloudera Navigator Encrypt](#) on page 2813 for new features in Cloudera Navigator Encrypt 3.13.0

Issues Fixed in Cloudera Navigator Encrypt 3.12.2

Cryptsetup Password Quality Check Skipped

The Navigator Encrypt `prepare` operation was failing on systems running RHEL 7.x because the 'Password quality check' of `cryptsetup` was enabled by default. It is possible that this could impact other operating systems.

The `navencrypt-prepare` command now skips the 'Password quality check' whenever it is enabled by default. This update has been implemented for all distributions supported by Navigator Encrypt.

Build `makefile` Fix

This patch fixes the build `makefile` and supports for following kernels:

- 3.10.0-514.2.2.el7.x86_64
- 3.10.0-514.16.1.el7.x86_64
- 3.10.0-327.el7.x86_64
- 3.10.0-514.10.2.el7.x86_64

Issues Fixed in Cloudera Navigator Encrypt 3.12.0

See [What's New in Cloudera Navigator Encrypt 3.12.0](#) on page 2813 for new features in Cloudera Navigator Encrypt 3.12.0.

Issues Fixed in Cloudera Navigator Encrypt 3.11.0

See [What's New in Cloudera Navigator Encrypt 3.11.0](#) on page 2813 for new features in Cloudera Navigator Encrypt 3.11.0.

Issues Fixed in Cloudera Navigator Encrypt 3.10.0

See [What's New in Cloudera Navigator Encrypt 3.10.0](#) on page 2813 for new features in Cloudera Navigator Encrypt 3.10.0.

Issues Fixed in Cloudera Navigator Encrypt 3.9.0

The software-signing GPG key is missing from the repository tarball for Navigator Encrypt

The GPG key used to sign the Navigator Encrypt packages is missing from the repository tarball, and must be imported from `archive.gazzang.com`.

Cloudera Bug: KT-2207

Issues Fixed in Cloudera Navigator Encrypt 3.8.0

Encrypting or decrypting a file or folder with a space in the name generates extraneous message

Running `navencrypt-move` to encrypt or decrypt a file or folder with a space in the name generates the following message at the command line:

```
/usr/sbin/navencrypt-move: line 233: [: /root/encrypt: binary operator expected
```

This message can be safely ignored.

Cloudera Bug: KT-1335

Running `navencrypt-move` twice on the same file does not fail properly

Running `navencrypt-move` twice on the same file succeeds on files with long paths with a message similar to the following:

```
ln: creating symbolic link '/path/to/file/to/encrypt' to  
'/path/to/mountpoint/category/path/to/file/to/encrypt': File exists
```

Cloudera Bug: KT-1707

Running `navencrypt-prepare --undo` removes all similar entries from `/etc/navencrypt/ztab`

Running `navencrypt-prepare --undo /path/to/mountpoint` removes all similar entries from `/etc/navencrypt/ztab`. For example, if you have mountpoints named `/mnt/data`, `/mnt/data1`, and `/mnt/data2`, running `navencrypt-prepare --undo /mnt/data` removes `/mnt/data`, `/mnt/data1`, and `/mnt/data2` from the Navigator Encrypt mount table (`/etc/navencrypt/ztab`).

Cloudera Bug: KT-1928

Encrypting or decrypting in parallel fails

Running `navencrypt-move` simultaneously in multiple terminals results in failure to encrypt or decrypt all of the specified data, with the following error message:

```
keyctl_unlink: Required key not available
```

No data is lost, because the source data is not removed, but you must re-run the failed operations sequentially.

Cloudera Bug: KT-1036

Restarting `navencrypt-mount` service after running `navencrypt-prepare` sometimes fails

Running `service navencrypt-mount restart` after `navencrypt-prepare` sometimes fails with the following error message:

```
WARNING: Another navencrypt-mount start process is currently running.
```

Cloudera Bug: KT-1206

Issues Fixed in Cloudera Navigator Encrypt 3.7.1

Pressing **Ctrl + C** during rule creation corrupts rule file

Pressing **Ctrl + C** before rule addition completes results in a corrupted rule file. Attempting to add further ACL rules fails with the following error:

```
[ERROR] Cannot parse ACL format: ACL header is not found: Did you type an incorrect key?
```

Intermittent Key Trustee Server communication errors

Intermittently, communication with the Key Trustee Server fails with an error similar to the following:

```
[error] UnicodeDecodeError: 'utf8' codec can't decode byte 0x80 in position 1: invalid start byte
```

The `navencrypt-move --per-file` option works only if the source and destination are on the same device

Using the `navencrypt-move --per-file` fails if the source and destination are on different devices.

Issues Fixed in Cloudera Navigator Encrypt 3.7.0

Kernel module reported by SLES as tainted kernel

The Navigator Encrypt kernel module is identified as a [tainted kernel](#) by SLES.

Navigator Encrypt does not always detect all processes accessing data to encrypt

Navigator Encrypt sometimes fails to detect all running processes that are accessing data that you want to encrypt.

Using `navencrypt-move` improperly fills storage device completely

Running `navencrypt-move` to encrypt a directory which contains the encryption mount point results in an infinite loop that fills the storage device.

Deprecated Items

The section lists OSs, Java versions, databases, platforms, CDH components and subcomponents, and product functionality that have been deprecated.

Terminology

Deprecated

Feature, component, platform, or functionality that Cloudera is planning to remove in a future release. Cloudera supports items that are deprecated until they are removed, and the deprecation gives customers time to plan for removal.

Removed

Feature, component, platform, or functionality that has been removed from the product and is no longer supported.

OSs, Java Versions, Databases, and Platforms

Item	Related Information	Release in Which Item Is Deprecated	Release in Which Item Is Removed
Operating System <ul style="list-style-type: none"> Ubuntu 14.04 	CDH and Cloudera Manager Supported Operating Systems on page 20	5.8.0	5.16.1, 6.0.0
Operating System <ul style="list-style-type: none"> SLES 11 	CDH and Cloudera Manager Supported Operating Systems on page 20	5.9.0	6.0.0
Operating System <ul style="list-style-type: none"> RHEL 5, CentOS 5, Oracle Enterprise Linux 5 Ubuntu 10.04 (already EOL by Canonical) and 12.04 Debian 7 	CDH and Cloudera Manager Supported Operating Systems on page 20	July 2015	6.0.0
Java 7	CDH and Cloudera Manager Supported JDK Versions on page 33	5.0.0	6.0.0
Database <ul style="list-style-type: none"> MySQL 5.0 PostgreSQL 8.1 		July 2015	6.0.0
Database <ul style="list-style-type: none"> MySQL 5.1 PostgreSQL 8.4 	CDH and Cloudera Manager Supported Databases on page 31	5.9.0	6.0.0
Database <ul style="list-style-type: none"> Oracle 11g 		5.9.0	6.0.0
Filesystem	Storing HBase Snapshots on Amazon S3 and Copying Cluster Data Using DistCp	5.5.0	5.7.0

Item	Related Information	Release in Which Item Is Deprecated	Release in Which Item Is Removed
Amazon S3 and S3n connectors. S3 and S3n are replaced by S3a.			
Tarball CDH Tarball Distribution		5.9.0	6.0.0
Cloudera Manager Tarball Distribution		5.9.0	6.0.0

CDH Components, Subcomponents, and Product Functionality

Item	Related Information	Release in Which Item Is Deprecated	Release in Which Item Is Removed
Activity Monitor	Activity Monitor is only used and deployed by Cloudera Manager when the MapReduce service (MRv1) is deployed.	5.9.0	7.0.0
DataFu		5.9.0	6.0.0
HFTP	Use WebHDFS	5.10.1	6.0.0
Hive CLI	About Hive	5.0.0	To be determined
HiveServer1	Starting HiveServer1 and the Hive Console in CDH	5.3.0	6.0.0
Legacy Scala clients for Kafka (consumer and producer)	The legacy Scala clients (producer and consumer) that are under the kafka-producer.* and kafka-consumer.* package.	CDK 3.0.0	6.1.0
Llama		5.5.0	6.0.0
MRv1, MapReduce v1 APIs, MapReduce service	Migrating from MapReduce (MRv1) to MapReduce (MRv2)	5.0.0	6.0.0
Mahout	Setting Up Apache Mahout Using the Command Line	5.5.0	6.0.0
Management of Key Trustee Server without Cloudera Manager	Cloudera Navigator Key Trustee Server	5.9.0	6.0.0
MR Pipes		5.9.0	6.0.0
Navigator Encrypt Filesystem-Level Encryption Using eCryptfs	Filesystem-Level Encryption with eCryptfs	July 2015	6.0.0
Old NameNode UI		5.5.0	6.0.0
Oozie Hive Action	Oozie Hive 2 Action Extension	5.7.0	To be determined

Deprecated Items

Item	Related Information	Release in Which Item Is Deprecated	Release in Which Item Is Removed
Parquet libraries named <code>com.twitter.*</code>	Using Apache Parquet Data Files with CDH	5.8.0	6.0.0
Sentry policy files	Migrating from Sentry Policy Files to the Sentry Service	5.8.0	6.0.0
Spark Standalone	Managing Spark Standalone Using the Command Line	5.5.0	6.0.0
Spark 1.x		5.13	6.0.0
Sqoop2		5.9.0	6.0.0
Whirr	Setting Up Apache Whirr Using the Command Line	5.5.0	6.0.0
YARN Capacity Scheduler	Use FairScheduler	5.9.0	6.0.0

Cloudera Manager Components, Subcomponents, and Product Functionality

Item	Related Information	Release in Which Item Is Deprecated	Release in Which Item is Removed
Multi Cloudera Manager Dashboard	Monitoring Multiple CDH Deployments Using the Multi Cloudera Manager Dashboard	5.11.1	6.0
Cloudera Manager API Versions v1 - 5	Cloudera Manager API Documentation	5.13	6.0

Appendix: Apache License, Version 2.0

SPDX short identifier: Apache-2.0

Apache License
Version 2.0, January 2004
<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License.

Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License.

Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims

licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution.

You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

1. You must give any other recipients of the Work or Derivative Works a copy of this License; and
2. You must cause any modified files to carry prominent notices stating that You changed the files; and
3. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
4. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions.

Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions.

Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6. Trademarks.

This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty.

Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability.

In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability.

While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

```
Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

    http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License.
```