

Release Notes 3

# Hortonworks Data Platform

**Date of Publish:** 2018-09-21

<http://docs.hortonworks.com>

# Contents

<b>HDP 3.0.1 Release Notes.....</b>	<b>3</b>
Component Versions.....	3
New Features.....	3
Deprecation Notices.....	4
Terminology.....	4
Removed Components and Product Capabilities.....	4
Unsupported Features.....	4
Technical Preview Features.....	4
Upgrading to HDP 3.0.1.....	5
Before you begin.....	5
Upgrade.....	5
Behavioral Changes.....	5
Apache Patch Information.....	7
Accumulo.....	7
Atlas.....	7
Calcite.....	8
DataFu.....	8
Hadoop.....	8
HBase.....	13
Hive.....	16
Kafka.....	16
Knox.....	16
Oozie.....	16
Phoenix.....	20
Pig.....	20
Ranger.....	20
Spark.....	22
Livy.....	23
Sqoop.....	23
Storm.....	23
Tez.....	24
Zeppelin.....	24
ZooKeeper.....	24
Fixed Common Vulnerabilities and Exposures.....	25
Fixed Issues.....	26
Known Issues.....	33
Documentation Errata.....	44
Legal Information.....	45

## HDP 3.0.1 Release Notes

This document provides you with the latest information about the Hortonworks Data Platform (HDP) 3.0.1 release and its product documentation.

### Component Versions

This section lists the official Apache versions of all the HDP 3.0.1 components. All components listed here are official Apache releases of the most recent stable versions available.

The Hortonworks approach is to provide patches only when necessary, to ensure the interoperability of components. Unless you are explicitly directed by Hortonworks Support to take a patch update, each of the HDP components should remain at the following package version levels, to ensure a certified and supported copy of HDP 3.0.0.

Official Apache versions for HDP 3.0.0:

- Apache Accumulo 1.7.0
- Apache Atlas 1.0.0
- Apache Calcite 1.16.0
- Apache DataFu 1.3.0
- Apache Druid 0.12.1 (incubating)
- Apache Hadoop 3.1.1
- Apache HBase 2.0.0
- Apache Hive 3.1.0
- Apache Kafka 1.1.1
- Apache Knox 1.0.0
- Apache Livy 0.5.0
- Apache Oozie 4.3.1
- Apache Phoenix 5.0.0
- Apache Pig 0.16.0
- Apache Ranger 1.1.0
- Apache Spark 2.3.1
- Apache Sqoop 1.4.7
- Apache Storm 1.2.1
- Apache Tez 0.9.1
- Apache Zeppelin 0.8.0
- Apache ZooKeeper 3.4.6

Later versions of a few Apache components are sometimes bundled in the HDP distribution in addition to the versions listed here. In this case, these later versions are listed in the [Technical Previews](#) table and should not substitute for the Apache component versions of the list provided here when used in a production environment.

Additional component versions:

- Apache Superset 0.23.3

**Note:** For information on open source software licensing and notices, refer to the Licenses and Notices files included with the software install package.

### New Features

This section highlights new features in HDP 3.0.1.

There are no new features introduced in this release.

## Deprecation Notices

This section points out any technology from previous releases that have been deprecated, moved, or removed from this release. Use this section as a guide for your implementation plans.

### Terminology

Items in this section are designated as follows:

#### Deprecated

Technology that Hortonworks is removing in a future HDP release. Marking an item as deprecated gives you time to plan for removal in a future HDP release.

#### Moving

Technology that Hortonworks is moving from a future HDP release and is making available through an alternative Hortonworks offering or subscription. Marking an item as moving gives you time to plan for removal in a future HDP release and plan for the alternative Hortonworks offering or subscription for the technology.

#### Removed

Technology that Hortonworks has removed from HDP and is no longer available or supported as of this release. Take note of technology marked as removed since it can potentially affect your upgrade plans.

### Removed Components and Product Capabilities

No components are deprecated or removed in this HDP release.

Please contact Hortonworks Support or your Hortonworks Account Team if you have any questions.

### Unsupported Features

This section lists features that are not officially supported by HDP.

Although the features listed in the Technical Preview and Community Features sections exist within HDP, Hortonworks does not currently support them.

### Technical Preview Features

The following features are available within HDP 3.0.1 but are not ready for production deployment. Hortonworks encourages you to explore these technical preview features in non-production environments and provide feedback on your experiences through the [Hortonworks Community Forums](#).

**Table 1: Technical Previews**

Component	Feature
Spark	Native Atlas connector or hook for Spark to capture metadata and lineage

Component	Feature
YARN	Introduced in a previous release: <ul style="list-style-type: none"> <li>• Cloud Service Metrics for YARN (<a href="#">YARN-4488</a>)</li> <li>• Cloud Service Health for YARN</li> <li>• In-Place upgrade of long-running applications in YARN</li> </ul>
Superset	Apache Superset is a modern, enterprise-ready business intelligence web application

## Upgrading to HDP 3.0.1

Instructions to upgrade from a previous version of HDP.

If you already have HDP 2.6.x or HDP 3.0.0 installed, upgrading your cluster to HDP 3.0.1 means:

- Keeping the same configuration files you used for HDP 2.6.x or HDP 3.0.0.
- Keeping the same data and metadata in the same location you used for HDP 2.6.x or HDP 3.0.0.
- Installing any new components (added for the first time in HDP 2.6.x or HDP 3.0.0) side-by-side with existing components.

### Before you begin

Consider this before you start your upgrade.

- Ensure that you know whether you want to do a major (HDP 2.6.x to HDP 3.0.1) or a minor (HDP 3.0.0 to HDP 3.0.1) upgrade.
- Ensure that you know which HDP components you need to upgrade at your installation.
- Read the Prepare to Upgrade section in the *Ambari Upgrade Guide*.

### Related Tasks

[Upgrade](#)

## Upgrade

Upgrade information and instructions.

You can upgrade your cluster using Ambari, see the *Ambari Upgrade Guide* for instructions.

### Related Concepts

[Before you begin](#)

### Related Information

[Ambari Major Upgrade Guide](#)

[Ambari Minor Upgrade Guide](#)

[Ambari Major Upgrade Guide for IBM Power Systems](#)

[Obtain the HDP Repositories](#)

[Register and Install HDP Version](#)

## Behavioral Changes

Behavioral changes denote a marked change in behavior from the previously released version to this version of HDP.

**Table 2: Behavioral Changes**

Hortonworks Bug ID	Apache Component	Apache JIRA	Summary	Details
BUG-109803	Ranger	N/A	Ranger authorization for CREATE TEMPORARY UDF operation in Hive.	<p>Scenario: Authorize temporary User Defined Function (UDF). Currently, the CREATE TEMPORARY UDF has a workaround solution of having a policy with Database=* and UDF= specified since the temp UDF is not associated to any database. From this release, UDF related operations are associated to a resource "Global". Hence users are expected to create policy to authorize Temp UDF instead of previous workaround.</p> <p>Previous behavior: Users need to create a policy with Database=* and UDF= specified.</p> <p>New behavior: UDF related operations are categorized to a resource "Global" and "Temp UDF admin" is added as a new permission, which authorizes ADD, COMPILE and CREATE TEMPORARY UDF. Hence previous workaround to have a "*" policy for DB and UDF to authorize temporary UDF commands is not required and will not work. Instead, users are expected to create this new policy with "Global" resource and "Temp UDF admin" permission for the authorization of Temp UDF. This permission should be treated carefully and only be given a few trusted users similar to Permanent UDF policy. Permanent UDFs are not affected by this change and continue to be handled by the existing UDF resource policy as before.</p> <p>Workaround/Expected Customer Action: If you have a previous workaround to authorize Hive Temp UDF, you must stop using the workaround and create new policy for "Global" resource.</p>
BUG-106681	Oozie	<a href="#">OOZIE-3328</a>	Oozie can understand Hive action once again in this release. Hive action is not supported on HDP# 3.0.0 but has backward compatibility in HDP# 3.0.1. The compatibility mode uses Beeline to attach the Hive server instead of using HiveCLI as earlier.	<p>Scenario: When I submit a workflow into Oozie, I can use my old hive actions and Oozie understands that.</p> <p>Previous behavior: Using Hive action was not possible in HDP# 3.0.0.</p> <p>New behavior: Using hive action is possible (so all the workflows which are created for HDP# 2.x) but highly recommend to rewrite workflows to switch to hive2 actions.</p> <p>Limitations: Limitations: On a wire encrypted cluster, Hive server certificates have to be spread on all the YARN nodes. Certificates should be imported into the Java's cacerts keystore.</p>

Hortonworks Bug ID	Apache Component	Apache JIRA	Summary	Details
BUG-109145	Zeppelin	<a href="#">ZEPELIN-3587</a>	Zeppelin is not showing complete output after running the paragraph.	<p>Scenario: User want to see the middle output of zeppelin paragraph.</p> <p>Previous behavior: Every output of each statement will be displayed, because zeppelin split the whole code into lines.</p> <p>New behavior: Only the last statement's output is displayed.</p> <p>Workaround/Expected Customer Action: If you want to display the output, you need to explicitly add a print statement.</p>

## Apache Patch Information

The following sections list patches in each HDP 3.0.1 component beyond what was fixed in the base version of the Apache component.

### Accumulo

This release provides Accumulo 1.7.0 with no additional Apache patches.

HDP 3.0.0 provided Accumulo 1.7.0 with no additional Apache patches.

### Atlas

This release provides Atlas 1.0.0 and the listed Apache patches.

- [ATLAS-2438](#): fix errors found in sitecheck - <https://whimsy.apache.org/site/project/atlas>
- [ATLAS-2523](#): Add HomeId and allow GUIDs to be specified on creates
- [ATLAS-2736](#): doc updated to replace reference to github with a link to download released artifacts
- [ATLAS-2745](#): AtlasEnumDefStore should remember default value
- [ATLAS-2785](#): Import Hive script should handle table name with database in -t option
- [ATLAS-2786](#): Honor ZK server ensemble provided in configuration
- [ATLAS-2788](#): Add policy permissions for tagsync user.
- [ATLAS-2793](#): updated AtlasClient approach of determining if kerberos authentication is enabled or not
- [ATLAS-2794](#): updated AtlasClient to explicitly specify JacksonJaxbJsonProvider
- [ATLAS-2801](#): Create a separate tar file for atlas-server components excluding atlas hook related components
- [ATLAS-2807](#): Re-evaluate classification propagation during entity delete
- [ATLAS-2808](#): Add proxy support
- [ATLAS-2809](#): updated authorization model to authorize add/update/remove of relationships
- [ATLAS-2812](#): Upgrade commons-fileupload to 1.3.3
- [ATLAS-2815](#): API, UI updates for faster rendering of entity details
- [ATLAS-2816](#): Allow ignoring relationship in EntityGraphRetriever for FullTextMapperV2
- [ATLAS-2817](#): Handle auto upgrade from janus 0.2.0 to janus 0.3.0
- [ATLAS-2817](#): Update Janusgraph artifact id
- [ATLAS-2817](#): JanusGraph 0.3.0
- [ATLAS-2819](#): UI: Re-evaluate classification propagation during entity delete
- [ATLAS-2821](#): Update old atlas-simple-authz-policy.json file with relationship permission attributes
- [ATLAS-2822](#): Provide option whether to delete propagated classification on entity delete during add classification
- [ATLAS-2823](#): updated hooks to support asynchronous notifications
- [ATLAS-2825](#): UI: Provide option whether to delete propagated classification on entity delete during add classification

- [ATLAS-2826](#): fix for failure in authorizing save-search operation.
- [ATLAS-2827](#): Indexed string compaction when length exceeds 33482223
- [ATLAS-2830](#): Tag Propagation, Entity Delete : Classification update notification not sent to entities down the lineage present after the deleted entity
- [ATLAS-2834](#): Init authorizer when Atlas is in HA mode
- [ATLAS-2835](#): Regression: Classifications propagated from term is not removed during term-entity disassociation
- [ATLAS-2836](#): Remove redundant encoding of vertex property keys [ATLAS-2838](#): Replace GraphHelper.encodeProperty() with AtlasGraphUtilsV2.encodeProperty()
- [ATLAS-2840](#): Regression - Tag Propagation, Deleted Entity : When an entity is deleted, tags propagated to it are removed

HDP 3.0.0 provided Atlas 1.0.0 and the following Apache patches:

- [ATLAS-2740](#): Updated import-hbase.sh script to populate values for missed HBase column-family attributes.
- [ATLAS-2751](#): Atlas is not consuming messages from ATLAS\_HOOK topic after recovering from zookeeper connection timeout.

## Calcite

This release provides Calcite 1.16.0 with no additional Apache patches.

HDP 3.0.0 provided Calcite 1.16.0 and no additional Apache patches.

## DataFu

This release provides DataFu 1.3.0 with no additional Apache patches.

HDP 3.0.0 provided DataFu 1.3.0 and no additional Apache patches.

## Hadoop

This release provides Hadoop Common 3.1.1 and no additional Apache patches.

HDP 3.0.0 provided Hadoop Common 3.1.0 and the following Apache patches:

- [HADOOP-15571](#): Multiple FileContexts created with the same configuration object should be allowed to have different umask. Contributed by Vinod Kumar Vavilapalli.
- [HDFS-13712](#): BlockReaderRemote.read() logging improvement. Contributed by Gergo Repas.
- [HDFS-13715](#): diskbalancer does not work if one of the blockpools are empty on a Federated cluster. Contributed by Bharat Viswanadham
- [HDFS-13715](#): Fix TestNameNodePrunesMissingStorages test failures on Windows. Contributed by Lukas Majercak.
- [HDFS-13635](#): Incorrect message when block is not found. Contributed by Gabor Bota.
- [HADOOP-15554](#): Improve JIT performance for Configuration parsing. Contributed by Todd Lipcon.
- [HDFS-13702](#): Remove HTrace hooks from DFSClient to reduce CPU usage. Contributed by Todd Lipcon.
- [HDFS-13703](#): Avoid allocation of CorruptedBlocks hashmap when no corrupted blocks are hit. Contributed by Todd Lipcon.
- [HADOOP-15574](#): Suppress build error if there are no docs after excluding private annotations. Contributed by Takanobu Asanuma.
- [HADOOP-15548](#): Randomize local dirs. Contributed by Jim Brennan.
- [HDFS-13705](#): The native ISA-L library loading failure should be made warning rather than an error message. Contributed by Shashikant Banerjee.
- [HADOOP-15458](#): TestLocalFileSystem#testFSOutputStreamBuilder fails on Windows. Contributed by Xiao Liang.
- [HDFS-13692](#): StorageInfoDefragmenter floods log when compacting StorageInfo TreeSet. Contributed by Bharat Viswanadham.
- [HDFS-13682](#): Cannot create encryption zone after KMS auth token expires.



- [HADOOP-15527](#): Improve delay check for stopping processes. Fixed script location. Contributed by Vinod Kumar Vavilapalli
- [HADOOP-15527](#): Improve delay check for stopping processes. Contributed by Vinod Kumar Vavilapalli
- [HADOOP-15533](#): Make WASB listStatus messages consistent. Contributed by Esfandiar Manii
- [HADOOP-15523](#): Shell command timeout given is in seconds whereas it is taken as millisecond while scheduling. Contributed by Bilwa S T.
- [HDFS-13681](#): Fix TestStartup.testNNFailToStartOnReadOnlyNNDir test failure on Windows. Contributed by Xiao Liang.
- [HDFS-13686](#): Add overall metrics for FSNamesystemLock. Contributed by Lukas Majercak.
- [HDFS-13174](#): hdfs mover -p /path times out after 20 min. Contributed by Istvan Fajth.
- [HDFS-13676](#): TestEditLogRace fails on Windows. Contributed by Zuoming Zhang.
- [HDFS-13673](#): TestNameNodeMetrics fails on Windows. Contributed by Zuoming Zhang.
- [HDFS-13675](#): Speed up TestDFSAdminWithHA. Contributed by Lukas Majercak.
- [HDFS-13563](#): TestDFSAdminWithHA times out on Windows. Contributed by Lukas Majercak.
- [HDFS-13641](#): Add metrics for edit log tailing. Contributed by Chao Sun.
- [BUG-103209](#): Dependency upgrades related to blackduck scan.
- [HADOOP-15527](#): Improve delay check for stopping processes. Contributed by Vinod Kumar Vavilapalli
- [HADOOP-15532](#): TestBasicDiskValidator fails with NoSuchFileException. Contributed by Giovanni Matteo Fumarola.
- [HDFS-13511](#): Provide specialized exception when block length cannot be obtained. Contributed by Gabor Bota.
- [HADOOP-15529](#): ContainerLaunch#testInvalidEnvVariableSubstitutionType is not supported in Windows. Contributed by Giovanni Matteo Fumarola.
- [HADOOP-15506](#): Upgrade Azure Storage Sdk version to 7.0.0 and update corresponding code blocks. Contributed by Esfandiar Manii.
- [HDFS-13653](#): Make dfs.client.failover.random.order a per nameservice configuration. Contributed by Ekanth Sethuramalingam.
- [HADOOP-15499](#): Performance severe drops when running RawErasureCoderBenchmark with NativeRSRawErasureCoder. Contributed by Sammi Chen.
- [HDFS-13667](#): Typo: Marking all datadoes as stale. Contributed by Nanda Kumar
- [HDFS-12670](#): can't renew HDFS tokens with only the hdfs client jar. Contributed by Arpit Agarwal.
- [HDFS-13664](#): Refactor ConfiguredFailoverProxyProvider to make inheritance easier. Contributed by Chao Sun.
- [HADOOP-15516](#): Add test cases to cover FileUtil#readLink. Contributed by Giovanni Matteo Fumarola.
- [HDFS-13659](#): Add more test coverage for contentSummary for snapshottable path. Contributed by Wei-Chiu Chuang.
- [HDFS-12950](#): [oiv] ls will fail in secure cluster. Contributed by Wei-Chiu Chuang.
- [HADOOP-15217](#): FsURLConnection does not handle paths with spaces. Contributed by Joseph Fourny and Zsolt Venczel.
- [HDFS-13545](#): "guarded" is misspelled as "gaured" in FSPermissionChecker.java. Contributed by Jianchao Jia.
- [HDFS-13650](#): Randomize baseDir for MiniDFSCluster in TestDFSStripedInputStream and TestDFSStripedInputStreamWithRandomECPolicy. Contributed by Anbang Hu.
- [HDFS-13649](#): Randomize baseDir for MiniDFSCluster in TestReconstructStripedFile and TestReconstructStripedFileWithRandomECPolicy. Contributed by Anbang Hu.
- [HDFS-13652](#): Randomize baseDir for MiniDFSCluster in TestBlockScanner. Contributed by Anbang Hu.
- [HADOOP-15137](#): ClassNotFoundException: org.apache.hadoop.yarn.server.api.DistributedSchedulingAMProtocol when using hadoop-client-minicluster. Contributed by Bharat Viswanadham
- [HDFS-13281](#): Namenode#createFile should be /.reserved/raw/ aware.. Contributed by Rushabh S Shah
- [HDFS-13155](#): BlockPlacementPolicyDefault.chooseTargetInOrder Not Checking Return Value for NULL. Contributed by Zsolt Venczel.
- [HDFS-13339](#): Volume reference can't be released and may lead to deadlock when DataXceiver does a check volume. Contributed by liaoyuxiangqin and Zsolt Venczel.

- [HDFS-13648](#): Fix TestGetConf#testGetJournalNodes on Windows due to a mismatch line separator. Contributed by Giovanni Matteo Fumarola.
- [HDFS-13651](#): TestReencryptionHandler fails on Windows. Contributed by Anbang Hu.
- MapReduce JHS should honor yarn.webapp.filter-entity-list-by-user(addendum). Contributed by Sunil Govindan.
- [HDFS-13637](#): RBF: Router fails when threadIndex (in ConnectionPool) wraps around Integer.MIN\_VALUE. Contributed by CR Hota.
- [HDFS-13636](#): Cross-Site Scripting vulnerability in HttpServer2 (Contributed by Haibo Yan via Daniel Templeton)
- [HDFS-13582](#): Improve backward compatibility for HDFS-13176 (WebHdfs file path gets truncated when having semicolon (;) inside). Contributed by Zsolt Venczel.
- [HDFS-13602](#): Add checkOperation(WRITE) checks in FSNamesystem. Contributed by Chao Sun.
- [HDFS-12978](#): Fine-grained locking while consuming journal stream. Contributed by Konstantin Shvachko. (cherry picked from commit 442dd87dcd3e9913a1d97793fde879ce51d701e7)
- [HDFS-13631](#): TestDFSAdmin#testCheckNumOfBlocksInReportCommand should use a separate MiniDFSCluster path. Contributed by Anbang Hu.
- BUG-104243:Dependency jackson-databind dependency related to blackduck scan.
- [HADOOP-15406](#): hadoop-nfs dependencies for mockito and junit are not test scope
- [HDFS-13611](#): Unsafe use of Text as a ConcurrentHashMap key in PBHelperClient.
- [HDFS-13629](#): Some tests in TestDiskBalancerCommand fail on Windows due to MiniDFSCluster path conflict and improper path usage. Contributed by Anbang Hu.
- [HDFS-13632](#): Randomize baseDir for MiniJournalCluster in MiniQJMHACluster for TestDFSAdminWithHA. Contributed by Anbang Hu.
- Additional check when unpacking archives. Contributed by Wilfred Spiegelenburg.
- [HADOOP-15497](#): TestTrash should use proper test path to avoid failing on Windows. Contributed by Anbang Hu.
- [HADOOP-15498](#): TestHadoopArchiveLogs (#testGenerateScript, #testPrepareWorkingDir) fails on Windows. Contributed by Anbang Hu.
- [HDFS-13591](#): TestDFSShell#testSetrepLow fails on Windows. Contributed by Anbang Hu.
- [HDFS-13627](#): TestErasureCodingExerciseAPIs fails on Windows. Contributed by Anbang Hu.
- [HADOOP-15449](#): Increase default timeout of ZK session to avoid frequent NameNode failover
- [HDFS-13628](#): Update Archival Storage doc for Provided Storage
- [HDFS-13620](#): Randomize the test directory path for TestHDFSFileSystemContract. Contributed by Anbang Hu.
- [HDFS-13619](#): TestAuditLoggerWithCommands fails on Windows. Contributed by Anbang Hu.
- [HADOOP-15473](#): Configure serialFilter in KeyProvider to avoid UnrecoverableKeyException caused by JDK-8189997. Contributed by Gabor Bota.
- [HDFS-13618](#): Fix TestDataNodeFaultInjector test failures on Windows. Contributed by Xiao Liang.
- [HADOOP-15486](#): Make NetworkTopology#netLock fair. Contributed by Nanda kumar.
- Skip the proxy user check if the ugi has not been initialized. Contributed by Daryn Sharp
- [HDFS-13540](#): DFSStripedInputStream should only allocate new buffers when reading. Contributed by Xiao Chen.
- [HADOOP-15478](#): WASB: hflush() and hsync() regression. Contributed by Thomas Marquardt.
- [HDFS-13493](#): Reduce the HttpServer2 thread count on DataNodes. Contributed by Erik Krogen.
- Additional check when unpacking archives. Contributed by Jason Lowe and Akira Ajisaka.
- [HDFS-13587](#): TestQuorumJournalManager fails on Windows. Contributed by Anbang Hu.
- [HDFS-13588](#): Fix TestFsDatasetImpl test failures on Windows. Contributed by Xiao Liang.
- [HADOOP-15450](#): Avoid fsync storm triggered by DiskChecker and handle disk full situation. Contributed by Arpit Agarwal.
- [HDFS-13601](#): Optimize ByteString conversions in PBHelper.
- [HDFS-13589](#): Add dfsAdmin command to query if upgrade is finalized. Contributed by Hanisha Koneru
- [HADOOP-15299](#): Bump Jackson 2 version to Jackson 2.9.x.
- [HDFS-13593](#): TestBlockReaderLocalLegacy#testBlockReaderLocalLegacyWithAppend fails on Windows. Contributed by Anbang Hu.

- [HDFS-13592](#): TestNameNodePrunesMissingStorages#testNameNodePrunesUnreportedStorages does not shut down cluster properly. Contributed by Anbang Hu.
- [HDFS-13586](#): Fsync fails on directories on Windows. Contributed by Lukas Majercack.
- [HDFS-13560](#): Insufficient system resources exist to complete the requested service for some tests on Windows. Contributed by Anbang Hu.
- [HDFS-13556](#): TestNestedEncryptionZones does not shut down cluster. Contributed by Anbang Hu.
- [HDFS-13554](#): TestDatanodeRegistration#testForcedRegistration does not shut down cluster. Contributed by Anbang Hu.
- [HDFS-13558](#): TestDatanodeHttpXFrame does not shut down cluster. Contributed by Anbang Hu.
- [BUG-103397](#). [HDFS-13589](#). Add dfsAdmin command to query if upgrade is finalized.
- [HADOOP-15469](#): S3A directory committer commit job fails if \_temporary directory created under dest. Contributed by Steve Loughran.
- [HADOOP-15250](#): Split-DNS MultiHomed Server Network Cluster Network IPC Client Bind Addr Wrong Contributed by Ajay Kumar
- [HDFS-13570](#): TestQuotaByStorageType,TestQuota,TestDFSOutputStream fail on Windows. Contributed by Anbang Hu.
- [HDFS-13581](#): DN UI logs link is broken when https is enabled. Contributed by Shashikant Banerjee.
- [HDFS-13559](#): TestBlockScanner does not close TestContext properly. Contributed by Anbang Hu.
- [HDFS-13512](#): WebHdfs getFileStatus doesn't return ecPolicy. Contributed by Ajay Kumar.
- [HDFS-13550](#):TestDebugAdmin#testComputeMetaCommand fails on Windows. Contributed by Anbang Hu.
- [HDFS-13557](#): TestDFSAdmin#testListOpenFiles fails on Windows. Contributed by Anbang Hu.
- [BUG-103055](#) Remove aiyun, openstack from hadoop-cloud-project
- [HDFS-13567](#): TestNameNodeMetrics#testGenerateEDEKTime,TestNameNodeMetrics#testResourceCheck should use a different cluster basedir. Contributed by Anbang Hu.
- [HDFS-13548](#): TestResolveHdfsSymlink#testFcResolveAfs fails on Windows. Contributed by Anbang Hu.
- [HDFS-11700](#). TestHDFSServerPorts#testBackupNodePorts doesn't pass on Windows. Contributed by Anbang Hu.
- [HDFS-13551](#): TestMiniDFSCluster#testClusterSetStorageCapacity does not shut down cluster. Contributed by Anbang Hu.
- [HADOOP-15442](#): ITestS3AMetrics.testMetricsRegister can't know metrics source's name. Contributed by Sean Mackrory.
- [HDFS-13544](#): Improve logging for JournalNode in federated cluster.
- [HDFS-13539](#) DFSStripedInputStream NPE when reportChecksumFailure.
- Add 2.9.1 release notes and changes documents
- [BUG-102169](#) Include Google Connector in Hadoop
- [HADOOP-15441](#): Log kms url and token service at debug level. Contributed by Gabor Bota
- [HDFS-13542](#): TestBlockManager#testNeededReplicationWhileAppending fails due to improper cluster shutdown in TestBlockManager#testBlockManagerMachinesArray on Windows. Contributed by Anbang Hu.
- [HDFS-13346](#): RBF: Fix synchronization of router quota and nameservice quota.
- [HADOOP-15454](#): TestRollingFileSystemSinkWithLocal fails on Windows. Contributed by Xiao Liang.
- [HDFS-13537](#): TestHdfsHelper does not generate jceks path properly for relative path in Windows. Contributed by Xiao Liang.
- [HDFS-12981](#): renameSnapshot a Non-Existent snapshot to itself should throw error. Contributed by Kittu Nanasi.
- [HADOOP-15446](#): WASB: PageBlobInputStream.skip breaks HBASE replication. Contributed by Thomas Marquardt
- [HADOOP-14841](#) Kms client should disconnect if unable to get output stream from connection. Contributed by Rushabh S Shah
- [HDFS-13525](#): RBF: Add unit test TestStateStoreDisabledNameservice. Contributed by Yiqun Lin.
- [HDFS-13481](#): TestRollingFileSystemSinkWithHdfs#testFlushThread: test failed intermittently (Contributed by Gabor Bota via Daniel Templeton)
- [HDFS-13512](#) . WebHdfs getFileStatus doesn't return ecPolicy. Contributed by Ajay Kumar.
- [HDFS-13488](#): RBF: Reject requests when a Router is overloaded. Contributed by Inigo Goiri.

- [HDFS-13503](#): Fix TestFsck test failures on Windows. Contributed by Xiao Liang.
- [HADOOP-15390](#): Yarn RM logs flooded by DelegationTokenRenewer trying to renew KMS tokens (xiaochen via rkanter)
- [HDFS-13283](#): Percentage based Reserved Space Calculation for DataNode. Contributed by Lukas Majercak.
- [HDFS-13509](#): Bug fix for breakHardlinks() of ReplicaInfo/LocalReplica, and fix TestFileAppend failures on Windows. Contributed by Xiao Liang.
- [HDFS-13508](#): RBF: Normalize paths (automatically) when adding, updating, removing or listing mount table entries. Contributed by Ekanth S.
- [HDFS-13484](#): RBF: Disable Nameservices from the federation. Contributed by Inigo Goiri.
- [HDFS-13434](#): RBF: Fix dead links in RBF document. Contributed by Chetna Chaudhari.
- [HDFS-13490](#): RBF: Fix setSafeMode in the Router. Contributed by Inigo Goiri.
- [HDFS-13499](#): RBF: Show disabled name services in the UI. Contributed by Inigo Goiri.
- [HDFS-13326](#): RBF: Improve the interfaces to modify and view mount tables. Contributed by Gang Li.
- [HADOOP-15411](#): AuthenticationFilter should use Configuration.getPropsWithPrefix instead of iterator. (Suma Shivaprasad via wangda)
- [HDFS-13336](#): Test cases of TestWriteToReplica failed in windows. Contributed by Xiao Liang.
- [HDFS-13356](#): Balancer:Set default value of minBlockSize to 10mb. Contributed by Bharat Viswanadham
- [BUG-101422](#): Disable HDFS and Hadoop unit tests that are flaky in Mawo.
- [HDFS-13408](#): MiniDFSCluster to support being built on randomized base directory. Contributed by Xiao Liang
- [HDFS-13433](#): webhdfs requests can be routed incorrectly in federated cluster. Contributed by Arpit Agarwal.
- [HDFS-13478](#): RBF: Disabled Nameservice store API. Contributed by Inigo Goiri.
- [HDFS-10183](#): Prevent race condition during class initialization. Contributed by Pavel Avgustinov.
- [HDFS-13453](#): RBF: getMountPointDates should fetch latest subdir time/date when parent dir is not present but / parent/child dirs are present in mount table. Contributed by Dibyendu Karmakar.
- [HADOOP-15396](#): Some java source files are executable
- [HDFS-13464](#): Fix javadoc in FsVolumeList#handleVolumeFailures. Contributed by Shashikant Banerjee
- [HADOOP-15366](#): Add a helper shutdown routine in HadoopExecutor to ensure clean shutdown. Contributed by Shashikant Banerjee.
- [HDFS-13426](#): Fix javadoc in FsDatasetAsyncDiskService#removeVolume. Contributed by Shashikant Banerjee.
- [HDFS-13463](#): Fix javadoc in FsDatasetImpl#checkAndUpdate. Contributed by Shashikant Banerjee.
- [HADOOP-15391](#): Add missing css file in hadoop-aws, hadoop-aliyun, hadoop-azure and hadoop-azure-datalake modules.
- [HDFS-12828](#): OIV ReverseXML Processor fails with escaped characters
- [BUG-100895](#): Ignore UTs failing due to environmental issues.
- [HDFS-13439](#): Add test case for read block operation when it is moved. Contributed by Ajay Kumar.
- [HDFS-13435](#): RBF: Improve the error loggings for printing the stack trace.
- [HDFS-13330](#): ShortCircuitCache#fetchOrCreate never retries. Contributed by Gabor Bota.
- [HDFS-13311](#): RBF: TestRouterAdminCLI#testCreateInvalidEntry fails on Windows. Contributed by Inigo Goiri
- [HDFS-13056](#): Expose file-level composite CRCs in HDFS which are comparable across different instances/ layouts. Contributed by Dennis Huo.
- [HDFS-13466](#): RBF: Add more router-related information to the UI.
- [HDFS-13462](#): Add BIND\_HOST configuration for JournalNode's HTTP and RPC Servers. Contributed by Lukas Majercak.
- [HDFS-13129](#): Add a test for DfsAdmin refreshSuperUserGroupsConfiguration. Contributed by Mukul Kumar Singh
- [HDFS-13433](#): webhdfs requests can be routed incorrectly in federated cluster. Contributed by Arpit Agarwal.
- [HDFS-11900](#): Hedged reads thread pool creation not synchronized. Contributed by John Zhuge.
- [HDFS-11043](#): TestWebHdfsTimeouts fails. Contributed by Xiaoyu Yao and Chao Sun.
- [HDFS-13315](#): Add a test for the issue reported in HDFS-11481 which is fixed by HDFS-10997. Contributed by Yongjun Zhang.
- [HADOOP-15332](#): Fix typos in hadoop-aws markdown docs. Contributed by Gabor Bota.

- [HADOOP-15331](#): Fix a race condition causing parsing error of java.io.BufferedInputStream in class org.apache.hadoop.conf.Configuration. Contributed by Miklos Szegedi.
- [HADOOP-15330](#) Remove jdk1.7 profile from hadoop-annotations module
- [HADOOP-15062](#): TestCryptoStreamsWithOpensslAesCtrCryptoCodec fails on Debian 9. Contributed by Miklos Szegedi.
- [HADOOP-14667](#): Flexible Visual Studio support. Contributed by Allen Wittenauer
- [HDFS-13438](#): Fix javadoc in FsVolumeList#removeVolume. Contributed by Shashikant Banerjee.
- [HDFS-13418](#): NetworkTopology should be configurable when enable DFSNetworkTopology. Contributed by Tao Jie.
- [HADOOP-14970](#): MiniHadoopClusterManager doesn't respect lack of format option. Contributed by Erik Krogen
- [HADOOP-15379](#): Make IrqHandler.bind() public. Contributed by Ajay Kumar
- [HDFS-13436](#): Fix javadoc of package-info.java
- [HDFS-7101](#): Potential null dereference in DFSck#doWork(). Contributed by skrho and Ted Yu.
- [HDFS-13427](#): Fix the section titles of transparent encryption document.
- [HDFS-13386](#): RBF: Wrong date information in list file(-ls) result. Contributed by Dibyendu Karmakar.
- [HDFS-13428](#): RBF: Remove LinkedList From StateStoreFileImpl.java. Contributed by BELUGA BEHR.
- [HDFS-13045](#): RBF: Improve error message returned from subcluster. Contributed by Inigo Goiri.
- [HADOOP-15346](#) S3ARetryPolicy for 400/BadArgument to be "fail". Contributed by Steve Loughran.
- [HADOOP-15313](#): TestKMS should close providers.
- [HADOOP-15357](#): Configuration.getPropsWithPrefix no longer does variable substitution. Contributed by Jim Brennan
- [HDFS-13328](#): Abstract ReencryptionHandler recursive logic in separate class. Contributed by Surendra Singh Lilhore.
- [HADOOP-15376](#): Remove double semi colons on imports that make Clover fall over.
- [HDFS-13420](#): License header is displayed in ArchivalStorage/MemoryStorage html pages. Contributed by Akira Ajisaka.
- [HADOOP-15374](#): Add links of the new features of 3.1.0 to the top page
- [HDFS-13384](#): RBF: Improve timeout RPC call mechanism. Contributed by Inigo Goiri.
- [HDFS-13410](#): RBF: Support federation with no subclusters. Contributed by Inigo Goiri.
- [HDFS-13380](#): RBF: mv/rm fail after the directory exceeded the quota limit. Contributed by Yiqun Lin.
- [HDFS-13388](#): RequestHedgingProxyProvider calls multiple configured NNs all the time. Contributed by Jinglun.
- [HDFS-13402](#): RBF: Fix java doc for StateStoreFileSystemImpl. Contributed by Yiran Wu.
- Added CHANGES/RELEASES/Jdiff for 3.1.0 release
- [HDFS-13353](#): RBF: TestRouterWebHDFSContractCreate failed. Contributed by Takanobu Asanuma.
- [HDFS-13350](#): Negative legacy block ID will confuse Erasure Coding to be considered as striped block. (Contributed by Lei (Eddy) Xu).
- [HDFS-13364](#): RBF: Support NamenodeProtocol in the Router. Contributed by Inigo Goiri.
- [HADOOP-15355](#): TestCommonConfigurationFields is broken by HADOOP-15312. Contributed by LiXin Ge.
- [HADOOP-15317](#): Improve NetworkTopology chooseRandom's loop.
- [HDFS-13289](#): RBF: TestConnectionManager#testCleanup() test case need correction. Contributed by Dibyendu Karmakar.
- [HADOOP-15253](#): Should update maxQueueSize when refresh call queue. Contributed by Tao Jie.
- [HADOOP-14999](#): AliyunOSS: provide one asynchronous multi-part based uploading mechanism. Contributed by Genmao Yu.
- [HADOOP-12862](#): LDAP Group Mapping over SSL can not specify trust store. Contributed by Wei-Chiu Chuang and Konstantin Shvachko.
- [HDFS-13087](#): Snapshotted encryption zone information should be immutable. Contributed by LiXin Ge.
- [HADOOP-15342](#): Updating ADLS connector to use the current SDK version (2.2.7). Contributed by Atul Sikaria.

## HBase

This release provides HBase 2.0.0 and the listed Apache patches.

- [HBASE-20772](#): Controlled shutdown fills Master log with the disturbing message 'No matching procedure found for rit=OPEN, location=ZZZZ, table=YYYYYY, region=XXXX transition to CLOSED'
- [HBASE-21018](#): RS crashed because AsyncFS was unable to update HDFS data encryption key
- [HBASE-21007](#): Memory leak in HBase REST server
- [HBASE-20885](#): Removed entry for RPC quota from hbase:quota when RPC quota is removed
- [HBASE-20989](#): Minor, miscellaneous logging fixes
- [HBASE-20935](#): HStore.removeCompactedFiles should log in case it is unable to delete a file (Vishal Khandelwal)
- [HBASE-20893](#): Data loss if splitting region while ServerCrashProcedure executing ADDENDUM: Rather than rollback, just do region reopens.
- [HBASE-20856](#): PITA having to set WAL provider in two places
- [HBASE-21012](#): Revert the change of serializing TimeRangeTracker
- [HBASE-20729](#): B&R BackupLogCleaner must ignore ProcV2 WAL files
- [HBASE-19369](#): Switch to Builder Pattern In WAL
- [HBASE-20939](#): There will be race when we call suspendIfNotReady and then throw ProcedureSuspendedException
- [HBASE-20921](#): Possible NPE in ReopenTableRegionsProcedure
- [HBASE-20859](#): [Backup and incremental load could fail in secure clusters]
- [HBASE-20966](#): RestoreTool#getTableInfoPath should look for completed snapshot only
- [HBASE-20927](#): RSGroupAdminEndpoint doesn't handle clearing dead servers if they are not processed yet.
- [HBASE-105255](#): debug changed to warn
- [HBASE-20867](#): RS may get killed while master restarts
- [HBASE-20873](#): Update doc for Endpoint-based Export
- [HBASE-20565](#): ColumnRangeFilter combined with ColumnPaginationFilter can produce incorrect result
- [HBASE-20878](#): Data loss if merging regions while ServerCrashProcedure executing
- [HBASE-20893](#): Data loss if splitting region while ServerCrashProcedure executing
- [HBASE-20401](#): Make MAX\_WAIT and waitIfNotFinished in CleanerContext configurable (Contributed by Stephen Wu)
- [HBASE-20914](#): Trim Master memory usage
- [HBASE-20870](#): Wrong HBase root dir in ITBLL's Search Tool
- [HBASE-20903](#): Backport HBASE-20792 "info:servername and info:sn inconsistent for OPEN region" to branch-2.0
- [HBASE-20869](#): Endpoint-based Export use incorrect user to write to destination
- [HBASE-20853](#): Polish "Add defaults to Table Interface so Implementors don't have to"
- [HBASE-20907](#): Fix Intermittent failure on TestProcedurePriority
- [HBASE-20875](#): MemStoreLABImp::copyIntoCell uses 7% CPU when writing
- [HBASE-20887](#): Backport HBASE-20865 "CreateTableProcedure is stuck in retry loop in CREATE\_TABLE\_WRITE\_FS\_LAYOUT state"
- [HBASE-20882](#): Backport HBASE-20616 "TruncateTableProcedure is stuck in retry loop in TRUNCATE\_TABLE\_CREATE\_FS\_LAYOUT state" to branch-2.0
- [HBASE-20884](#): Replace Base64 with j.u.Base64
- [HBASE-19572](#): RegionMover should use the configured default port number and not the one from HConstants
- [HBASE-20697](#): Can't cache All region locations of the specify table by calling table.getRegionLocator().getAllRegionLocations()
- [HBASE-20860](#): Merged region's RIT state may not be cleaned after master restart
- [HBASE-20833](#): Modify pre-upgrade coprocessor validator to support table level coprocessors
- [HBASE-20806](#): Split style journal for flushes and compactions
- [HBASE-20854](#): Wrong retries number in RpcRetryingCaller's log message
- [HBASE-20822](#): TestAsyncNonMetaRegionLocator is flakey
- [HBASE-20691](#) : Change the default WAL storage policy back to "NONE"
- [HBASE-20474](#): Show non-RPC tasks on master/regionserver Web UI by default

- [HBASE-20839](#): Fallback to FSHLog if we can not instantiated AsyncFSWAL when user does not specify AsyncFSWAL explicitly
- [HBASE-20789](#): TestBucketCache#testCacheBlockNextBlockMetadataMissing is flaky
- [HBASE-20447](#): Only fail cacheBlock if block collisions aren't related to next block metadata
- [HBASE-20244](#): NoSuchMethodException when retrieving private method decryptEncryptedDataEncryptionKey from DFSCClient
- [HBASE-20826](#): Truncate really long RpcServer warnings unless TRACE is on
- [HBASE-20817](#): Infinite loop when executing ReopenTableRegionsProcedure
- [HBASE-20825](#): Fix pre and post hooks of CloneSnapshot and RestoreSnapshot for Access checks
- [HBASE-20812](#): Add defaults to Table Interface so implementors don't have to Signed-off-by: Mike Drob <mdrob@apache.org>
- [HBASE-20810](#) : Include the procedure id in the exception message in HBaseAdmin for better debugging
- [HBASE-19722](#): Meta query statistics metrics source
- [HBASE-20781](#): Save recalculating families in a WALEdit batch of Cells
- [HBASE-20785](#): NPE getting metrics in PE testing scans
- [HBASE-20769](#): getSplits() has a out of bounds problem in TableSnapshotInputFormatImpl
- [HBASE-20795](#): Allow option in BBKVComparator.compare to do comparison without sequence id
- [HBASE-20777](#): RpcConnection could still remain opened after we shutdown the NettyRpcServer
- [HBASE-20780](#) : ServerRpcConnection logging cleanup Get rid of one of the logging lines in ServerRpcConnection by amalgamating all into one new-style log line.
- [HBASE-20403](#): Fix race between prefetch task and non-read HFile reads
- [HBASE-20770](#): WAL cleaner logs way too much; gets clogged when lots of work to do
- [HBASE-19764](#): Fix Checkstyle errors in hbase-endpoint
- [HBASE-20778](#): Make it so WALPE runs on DFS
- [HBASE-20710](#): extra cloneFamily() in Mutation.add(Cell); ADDENDUM checkstyle
- [HBASE-20710](#): extra cloneFamily() in Mutation.add(Cell) Signed-off-by: Michael Stack <stack@apache.org>
- [HBASE-20752](#): Make sure the regions are truly reopened after ReopenTableRegionsProcedure
- [HBASE-20767](#): Always close hbaseAdmin along with connection in HBTU
- [HBASE-20642](#): Clients should re-use the same nonce across DDL operations
- [HBASE-20765](#): Move branch-2.0 version to 2.0.2-SNAPSHOT from 2.0.1
- [HBASE-20706](#): Prevent MTP from trying to reopen non-OPEN regions
- [HBASE-20745](#): Log when master proc wal rolls
- [HBASE-20681](#): Explicitly include hamcrest in binary tarball
- [HBASE-20733](#): QABot should run checkstyle tests if the checkstyle configs change
- [HBASE-20726](#): Update RELEASNOTES.md and CHANGES.md for 2.0.1RC
- [HBASE-20725](#): Copy down docs from master branch (absent chapters on features not in 2.0.x)
- [HBASE-20699](#):: QuotaCache should cancel the QuotaRefresherChore service inside its stop()
- [HBASE-20721](#):: Set version on branch-2.0 to 2.0.1 in prep for 2.0.1 RC

HDP 3.0.0 provided HBase 2.0.0 and the following Apache patches:

- [HBASE-20723](#): Custom hbase.wal.dir results in data loss because we write recovered edits into a different place than where the recovering region server looks for them
- [HBASE-20722](#): Make RegionServerTracker only depend on children changed event
- [HBASE-20668](#): Avoid permission change if ExportSnapshot's copy fails
- [HBASE-20657](#): Retrying RPC call for ModifyTableProcedure may get stuck
- [HBASE-20635](#): Add missing shaded-pb conversion method for UserPermission
- [HBASE-20627](#): Relocate RS Group pre/post hooks from RSGroupAdminServer to RSGroupAdminEndpoint
- [HBASE-20579](#): Include original exception in wrapped exception
- [HBASE-20579](#): Improve snapshot manifest copy in ExportSnapshot
- [HBASE-20547](#): Restore from backup will fail if done from a different file system
- [HBASE-20531](#): RS may throw NPE when close meta regions in shutdown procedure.

- [HBASE-20495](#): REST unit test fails with NoClassDefFoundError against hadoop3
- [HBASE-20196](#): Maintain all regions with same size in memstore flusher
- [HBASE-20136](#): TestKeyValue misses ClassRule and Category annotations
- [HBASE-19735](#): Create a client-tarball assembly
- [HBASE-19478](#): Utilize multi-get to speed up WAL file checking in BackupLogCleaner (Toshihiro Suzuki)
- [HBASE-18135](#): Implement mechanism for RegionServers to report file archival for space quotas
- [HBASE-18133](#): Decrease quota reaction latency by HBase
- [HBASE-17549](#): HBase-Spark Module: Corrected - Incorrect log at println and unwanted comment cod

## Hive

This release provides Hive 3.1.0 with no additional Apache patches.

HDP 3.0.0 provided Hive 3.1.0 and no additional Apache patches.

## Kafka

This release provides Kafka 1.1.1 with no additional Apache patches.

HDP 3.0.0 provided Kafka 1.0.1 and no additional Apache patches.

A new feature was added to capture producer and topic partition level metrics. This feature is available in HDP 3.0.0. It is not available in Apache Kafka 1.0.1.

## Knox

This release provides Knox 1.0.0 with no Apache patches.

HDP 3.0.0 provided Knox 1.0.0 and no additional Apache patches.

## Oozie

This release provides Oozie 4.3.1 and the listed Apache patches.

- [OOZIE-3304](#): Parsing sharelib timestamps is not threadsafe (dionusos, matijhs via andras.piros)
- [OOZIE-3135](#): Configure log4j2 in SqoopMain
- [OOZIE-3072](#): oozie.service.HadoopAccessorService.action.configurations should overwrite default values set in Hadoop's configuration files (gezapeti)
- [OOZIE-2916](#): Set a job name for the MR Action's child job (asasvari)
- [OOZIE-2494](#): Cron syntax not handling DST properly (kmarion via andras.piros)
- [OOZIE-2726](#): Flaky test due to daylight saving changes (sasishsaley, andras.piros via gezapeti)
- [OOZIE-3173](#) : Coordinator job with frequency using cron syntax creates only one action in catchup mode (andras.piros)
- [OOZIE-2691](#): Show workflow action retry information in UI
- [OOZIE-3109](#): [log-streaming] Escape HTML-specific characters (dionusos via andras.piros)
- [OOZIE-2815](#): Oozie not always display job log
- [OOZIE-3075](#): Follow-up on OOOZIE-3054: create the lib directory if it doesn't exist (pbacsko)
- [OOZIE-3054](#): Disable erasure coding for sharelib if Oozie runs on Hadoop 3 (pbacsko)
- [OOZIE-3048](#): Check El Functions for the coordinator action (satishsaley)
- [OOZIE-3004](#) : Forked action retry info is not working
- [OOZIE-2691](#): Show workflow action retry information in UI
- [OOZIE-3009](#): Number of Oozie tests executed dropped after OOOZIE-2854 (andras.piros via asasvari)
- [OOZIE-2854](#): Oozie should handle transient database problems (andras.piros via gezapeti)
- [OOZIE-3026](#): fix openjpa enhancer stage during build for logging (dbdist13, andras.piros via pbacsko)
- [OOZIE-2816](#): Strip out the first command word from Sqoop action if its "sqoop" (harsh)
- [OOZIE-3113](#): Retry for ZK lock release (satishsaley)
- [OOZIE-3208](#): "It should never happen" error messages should be more specific to root cause (kmarion via andras.piros)



- [OOZIE-2940](#): Possible NullPointerException in WorkflowActionBean (dionusos via gezapeti)
- [OOZIE-3297](#): Retry logic does not handle the exception from BulkJPAExecutor properly (andras.piros)
- [OOZIE-2733](#): change org.apache.hadoop.fs.permission.AccessControlException to org.apache.hadoop.security.AccessControlException (gezapeti)
- [OOZIE-3021](#) : Error on job or SLA event listening: WARN logs instead of DEBUG (andras.piros via asasvari)
- [OOZIE-2387](#): Oozie is Unable to handle Spaces in file/archive tag. (asasvari via gezapeti)
- [OOZIE-2818](#): Can't overwrite oozie.action.max.output.data on a per-workflow basis (asasvari via pbacsko)
- [OOZIE-2041](#): Add an admin command to run the PurgeXCommand (abhishekbafna)
- [OOZIE-2875](#): Typo in ssh action twiki docs (Dongying Jiao via gezapeti)
- [OOZIE-1401](#): PurgeCommand should purge the workflow jobs w/o end\_time (asasvari)
- [OOZIE-3156](#): Retry SSH action check when cannot connect to remote host (txsing via andras.piros)
- [OOZIE-3183](#) :Better logging for SshActionExecutor and extended HA capability when calling to remote host (andras.piros)
- [OOZIE-2769](#): Extend FS action to allow setrep on a file (Artem Ervits via gezapeti)
- [OOZIE-3085](#): Improve logging in ActionExecutors: add log entry for start(), end() and kill() (kmarion via andras.piros)

HDP 3.0.0 provided Oozie 4.3.1 and the following Apache patches:

- [OOZIE-2802](#):: Spark action failure on Spark 2.1.0 due to duplicate sharelibs.
- [OOZIE-2658](#):: --driver-class-path can overwrite the classpath in SparkMain.
- [OOZIE-2787](#):: Oozie distributes application jar twice making the spark job fail.
- [OOZIE-2606](#):: Set spark.yarn.jars to fix Spark 2.0 with Oozie.
- [OOZIE-2436](#):: Fork/join workflow fails with "oozie.action.yarn.tag must not be null".
- [OOZIE-2578](#):: Oozie example distcp job fails to run within an encrypted zone with checksum match error.
- [OOZIE-2362](#):: SQL injection in BulkJPAExecutor.
- [OOZIE-2577](#):: Flaky tests TestCoordActionInputCheckXCommand.testTimeout and testTimeoutWithException.
- [OOZIE-2570](#):: remove -PtestPatchCompile from patch testing as there is no such profile.
- [OOZIE-2504](#):: Create a log4j.properties under HADOOP\_CONF\_DIR in Shell Action.
- [OOZIE-2567](#):: HCat connection is not closed while getting heat cred.
- [OOZIE-2547](#):: Add mapreduce.job.cache.files to spark action.
- [OOZIE-2550](#):: Flaky tests in TestZKUUIDService.java.
- [OOZIE-2445](#):: Doc for - Specifying coordinator input datasets in more logical ways.
- [OOZIE-2541](#):: Possible resource leak in Hive2Credentials.
- [OOZIE-2563](#):: Pass spark-defaults.conf to spark action.
- [OOZIE-2556](#):: TestAbandonedCoordChecker.testCatchupJob is flaky.
- [OOZIE-2522](#):: There can be multiple coord submit from bundle in case of ZK glitch.
- [OOZIE-2553](#):: Cred tag is required for all actions in the workflow even if an action does not require it.
- [OOZIE-2503](#):: show ChildJobURLs to spark action.
- [OOZIE-2551](#):: Feature request: epoch timestamp generation.
- [OOZIE-2542](#):: Option to disable OpenJPA BrokerImpl finalization.
- [OOZIE-2447](#):: Illegal character 0x0 oozie client.
- [OOZIE-2548](#):: Flaky test TestZKLocksService.testLockReaper.
- [OOZIE-2546](#):: Improperly closed resources in OozieDBCLI.
- [OOZIE-2476](#):: When one of the action from fork fails with transient error, WF never joins .
- [OOZIE-2475](#):: Oozie does not cleanup action dir of killed actions.
- [OOZIE-2535](#):: User can't disable uber mode.
- [OOZIE-2482](#):: Pyspark job fails with Oozie.
- [OOZIE-2467](#):: Oozie can shutdown itself on long GC pause.
- [OOZIE-2537](#):: SqoopMain does not set up log4j properly.
- [OOZIE-2532](#):: patch apply does not handle binary files.

- [OOZIE-2330](#):: Spark action should take the global jobTracker and nameNode configs by default and allow file and archive elements.
- [OOZIE-2510](#):: TestLogStreamer.testStreamLog. fails in other timezones.
- [OOZIE-2531](#):: Prevent Spark trying for token which is already available.
- [OOZIE-2509](#):: SLA job status can stuck in running state.
- [OOZIE-2529](#):: Support adding secret keys to Credentials of Launcher.
- [OOZIE-1402](#):: Increase retry interval for non-progressing coordinator action with fix value .
- [OOZIE-2512](#):: ShareLibservice returns incorrect path for jar.
- [OOZIE-2471](#):: Show child job url tab for distcp.
- [OOZIE-2511](#):: SubWorkflow missing variable set from option if config-default is present in parent workflow.
- [OOZIE-2391](#):: spark-opts value in workflow.xml is not parsed properly.
- [OOZIE-2489](#):: XML parsing is vulnerable.
- [OOZIE-2485](#):: Oozie client keeps trying to use expired auth token.
- [OOZIE-2490](#):: Oozie can't set hadoop.security.token.service.use\_ip.
- [OOZIE-2474](#):: <job-xml> is not being applied to the launcher job.
- [OOZIE-2486](#):: TestSLAEventsGetForFilterJPAExecutor is flakey.
- [OOZIE-2481](#):: Add YARN\_CONF\_DIR in the Shell action.
- [OOZIE-2492](#):: JSON security issue in js code.
- [OOZIE-2429](#):: TestEventGeneration test is flakey.
- [OOZIE-2466](#):: Repeated failure of TestMetricsInstrumentation.testSamplers.
- [OOZIE-2470](#):: Remove infinite socket timeouts in the Oozie email action.
- [OOZIE-2246](#):: CoordinatorInputCheckCommand does not behave properly when har file is one of data dependency and doesn't exist.
- [OOZIE-2446](#):: Job does not fail during submission if non existent credential is specified .
- [OOZIE-2283](#):: Documentation should not say that System.exit is not allowed in Java Action .
- [OOZIE-2400](#):: Workflow xml configuration parser cannot deal with namespace prefix.
- [OOZIE-2452](#):: Coordinator Functional Specification - EL Constants Typo.
- [OOZIE-2173](#):: DISCLAIMER.txt still refers to Apache Incubator.
- [OOZIE-2312](#):: oozie doesn't purge audit and error log.
- [OOZIE-2431](#):: Support hive.metastore.uris as alias for hcat.metastore.uri.
- [OOZIE-1976](#):: Specifying coordinator input datasets in more logical ways.
- [OOZIE-2444](#):: Need conditional logic in bundles.
- [OOZIE-2394](#):: Oozie can execute command without holding lock.
- [OOZIE-1922](#):: MemoryLocksService fails if lock is acquired multiple times in same thread and released.
- [OOZIE-2432](#):: TestPurgeXCommand fails.
- [OOZIE-2434](#):: inconsistent coord action status and workflow job status.
- [OOZIE-2438](#):: Oozie client "jobs -filter" diagnostic message clarification.
- [OOZIE-2165](#):: Job log fetching can fail in Oozie HA mode when using doAs impersonation.
- [OOZIE-2430](#):: Add root logger for hive,sqoop action.
- [OOZIE-2441](#):: SubWorkflow action with propagate-configuration but no global section throws NPE on submit.
- [OOZIE-2370](#):: Docs for Coordinator Action Status Notification has wrong property name.
- [OOZIE-2419](#):: HBase credentials are not correctly proxied.
- [OOZIE-2439](#):: FS Action no longer uses name-node from global section or default NN.
- [OOZIE-2435](#):: TestCoordChangeXCommand is flakey.
- [OOZIE-2428](#):: TestSLAService, TestSLAEventGeneration flakey tests.
- [OOZIE-2437](#):: Exclude additional Hadoop dependencies from Spark Sharelib.
- [OOZIE-2030](#):: Configuration properties from global section is not getting set in Hadoop job conf when using sub-workflow action in Oozie workflow.xml.
- [OOZIE-2380](#):: Oozie Hive action failed with wrong tmp path.
- [OOZIE-2222](#):: Oozie UI parent job should be clickable.

- [OOZIE-2407](#):: AbandonedService should not send mail if there is no abandoned coord.
- [OOZIE-2402](#):: oozie-setup.sh sharelib create takes a long time on large clusters.
- [OOZIE-2185](#):: Make oozie cli source conf/oozie-client-env.sh.
- [OOZIE-2413](#):: Kerberos credentials can expire if the KDC is slow to respond.
- [OOZIE-2411](#):: Add BCC to oozie email action.
- [OOZIE-2397](#):: LAST\_ONLY and NONE don't properly handle READY actions.
- [OOZIE-2401](#):: Typo in twiki docs with FilesAchives instead of FilesArchives.
- [OOZIE-2168](#):: Oozie flow and action names have 50 char limit.
- [OOZIE-2346](#):: Add sub-workflow information like the super-parent id and workflow depth into the 'oozie.job.info' property.
- [OOZIE-2303](#):: Typo in documentation.
- [OOZIE-2328](#):: Coordinator endtime change should check if the last action is in database .
- [OOZIE-2367](#):: fs delete should support skipTrash option.
- [OOZIE-2368](#):: coord:dateOffset and coord:dateTzOffset can't be used for coord initial-instance.
- [OOZIE-2369](#):: coord:offset doesn't resolve correctly.
- [OOZIE-2251](#):: Expose instrumental matrices in Realtime Graphing tool.
- [OOZIE-2314](#):: Unable to kill old instance child job by workflow or coord rerun by Launcher .
- [OOZIE-2386](#):: org.apache.oozie.util.TestXConfiguration.testSubstituteVar is flakey.
- [OOZIE-2385](#):: org.apache.oozie.TestCoordinatorEngineStreamLog.testCoordLogStreaming is flakey.
- [OOZIE-2382](#):: org.apache.oozie.action.hadoop.TestPigMain.testPig\_withNullExternalID is flakey.
- [OOZIE-2379](#):: org.apache.oozie.command.coord.TestCoordPushDependencyCheckXCommand.testRequeueOnException is flakey.
- [OOZIE-2378](#):: org.apache.oozie.client.TestOozieCLI.testAdminInstrumentation is flakey.
- [OOZIE-2377](#):: Hive2 Action should not propagate oozie.hive2.\* properties to Beeline.
- [OOZIE-2376](#):: Default action configs not honored if no <configuration> section in workflow.
- [OOZIE-2365](#):: oozie fail to start when smtp password not set.
- [OOZIE-2360](#):: Spark Action fails due to missing mesos jar.
- [OOZIE-2345](#):: Parallel job submission for forked actions.
- [OOZIE-2358](#):: Coord rerun cleanup should reuse hcat connections.
- [OOZIE-2356](#):: Add a way to enable/disable credentials in a workflow.
- [OOZIE-2355](#):: Hive2 Action doesn't pass along oozie configs to jobconf.
- [OOZIE-2318](#):: Provide better solution for specifying SSL truststore to Oozie Client.
- [OOZIE-2344](#):: Enabling 'oozie.action.jobinfo.enable' doesn't inject the job information into the map/reduce job's configuration..
- [OOZIE-2350](#):: Package changes for release.
- [OOZIE-2347](#):: Remove unnecessary new Configuration./new jobConf. calls from oozie.
- [OOZIE-2348](#):: Recovery service keeps on recovering coord action of suspended jobs.
- [OOZIE-2277](#):: Honor oozie.action.sharelib.for.spark in Spark jobs.
- [OOZIE-2322](#):: Oozie Web UI doesn't work with Kerberos in Internet Explorer 10 or 11 and curl .
- [OOZIE-2343](#):: Shell Action should take Oozie Action config and setup HADOOP\_CONF\_DIR.
- [OOZIE-2245](#):: Service to periodically check database schema.
- [OOZIE-2332](#):: Add ability to provide Hive and Hive 2 Action queries inline in workflows .
- [OOZIE-2329](#):: Make handling yarn restarts configurable.
- [OOZIE-2228](#):: Statustransit service doesn't pick bundle with suspend status.
- [OOZIE-2325](#):: Shell action fails if user overrides oozie.launcher.mapreduce.map.env.
- [OOZIE-2324](#):: A syntax error in the kill node causes the workflow to get stuck and other problems.
- [OOZIE-2309](#):: Enable the coord:dateOffset. function in /coordinator-app/datasets/dataset/@initial-instance.
- [OOZIE-2305](#):: Compile Oozie with Hive-1.2.0.

- [OOZIE-2320](#):: TestZKXLogStreamingService.testStreamingWithMultipleOozieServers\_coordActionList is failing.
- [OOZIE-2293](#):: Oozie 4.0.1 build failed while building Catalog.
- [OOZIE-2308](#):: Add support for bundle:conf. function.
- [OOZIE-2315](#):: TestOozieCLI.testshareLibUpdate\_withSecurity fails with Hadoop 2.
- [OOZIE-2291](#):: Hive2 workflow.xml.security should have "cred" in action tag instead of "hive2" tag.
- [OOZIE-2289](#):: hive-jdbc dependency in core/pom.xml should be compile.
- [OOZIE-2300](#):: TestAuthFilterAuthOozieClient.testClientAuthTokenCache fails with Hadoop 2.7.0 and later.
- [OOZIE-1837](#):: LauncherMainHadoopUtils sensitive to clock skew.
- [OOZIE-2187](#):: Add a way to specify a default JT/RM and NN.
- [OOZIE-2272](#):: Use Hadoop's CredentialProvider for passwords in oozie-site.
- [OOZIE-2287](#):: Add support for deleting hcat partitions in fs action delete.
- [OOZIE-2285](#):: Change in concurrency should trigger coord action ready command.
- [OOZIE-2284](#):: HBaseCredentials should only add hbase-default.xml and hbase-site.xml to actionConf.
- [OOZIE-2286](#):: Update Log4j and Log4j-extras to latest 1.2.x release.
- [OOZIE-2250](#):: Show log for WAITING and TIMEDOUT coord actions.
- [OOZIE-2262](#):: Fix log streaming from other server with start/end filter.
- [OOZIE-2159](#):: 'oozie validate' command should be moved server-side.
- [OOZIE-2271](#):: Upgrade Tomcat to 6.0.44.
- [OOZIE-2266](#):: Fix 'total' actions returned in coordinator job info.
- [OOZIE-2264](#):: Fix coord:offset. to resolve correct data set.
- [OOZIE-2178](#):: fix javadoc to compile on JDK8.
- [OOZIE-2268](#):: Update ActiveMQ version for security and other fixes.
- [OOZIE-2215](#):: Support glob in FS EL function.
- [OOZIE-2270](#):: Audit log doesn't get loaded when "audit log" tab is clicked for coord job.
- [OOZIE-2257](#):: BundleStatusTransitXCommand doesn't update last modified time.
- [OOZIE-2260](#):: Only set "java.io.tmpdir" to "./tmp" for the AM.

## Phoenix

This release provides Phoenix 5.0.0 with the following Apache patches.

- [PHOENIX-4809](#): Only cache PhoenixConnections when lease renewal is on
- [PHOENIX-3991](#): ROW\_TIMESTAMP on TIMESTAMP column type throws ArrayOutOfBound when upserting without providing a value.
- [PHOENIX-4825](#): Replace usage of HBase Base64 implementation with java.util.Base64(related to HBASE-20884)

HDP 3.0.0 provided Phoenix 5.0.0 and no additional Apache patches.

## Pig

This release provides Pig 0.16.0 and the same list of patches listed for HDP 3.0.0.

HDP 3.0.0 provided Pig 0.16.0 and the following Apache patches:

- [PIG-4947](#): LOAD with HBaseStorage using a mix of pure wildcards and prefixed wildcards results in empty maps for the pure wildcarded column families.
- [PIG-4948](#) : Pig on Tez AM use too much memory on a small cluster.
- [PIG-4921](#): Kill running jobs on InterruptedException (partial).
- [PIG-4916](#): Pig on Tez fail to remove temporary HDFS files in some cases.
- [PIG-5287](#): Bump jython to 2.7.1.

## Ranger

This release provides Ranger 1.1.0 and the same list of patches listed for HDP 3.0.0.

HDP 3.0.0 provided Ranger 1.1.0 and the following Apache patches:

- [RANGER-2201](#): Fix to log no ranger audits when entityId value is not null or empty string
- [RANGER-2195](#): TagPolicy not working due to failure to update tag policy version --BUG-109598
- [RANGER-2188](#): Support multiple threads to build Trie and on-lookup post-setup for Trie nodes - set default to lazy post-setup
- [RANGER-2167](#): Upgrade to Apache parent pom version 20
- [RANGER-2191](#): Update ranger-tool with new options to control Trie
- [RANGER-2186](#): Increment service-specific policy and tag versions after update transaction is committed
- [RANGER-2189](#): default policy for Atlas service should allow relationship-type access to all
- [RANGER-2114](#): Internal Exception: com.mysql.jdbc.MySQLDataTruncation: Data truncation: Data too long for
- [RANGER-2138](#): Internal Exception: com.mysql.jdbc.MySQLDataTruncation: Data truncation: Data too long for
- [RANGER-2177](#): Handle validations for duplicate configuration item during service create/edit
- [RANGER-2162](#): Upgrade c3p0 libraries
- [RANGER-2168](#): dd service admin user through service config
- [RANGER-2169](#): Create unique index on service and name column of x\_policy tablet
- [RANGER-2153](#): Supply the function of reverting policy history version.
- [RANGER-2184](#): Update RangerAtlas authorization to authorize add/update/remove of relationships
- [RANGER-2180](#): Handle token replacement correctly when token is not defined in the request context
- [RANGER-2177](#): Handle validations for duplicate configuration item during service create/edit
- [RANGER-2152](#): Incorrect debugging information in RangerPluginClassLoader.java
- [RANGER-2150](#): Unit test coverage for XUserMgr and UserMgr class
- [RANGER-2158](#): Performance improvement to REST API call to update policy

HDP 3.0.0 provided Ranger 1.1.0 and the following Apache patches:

- [RANGER-1715](#): Enhance Ranger Hive Plugin to support authorization on Hive replication Tasks
- [RANGER-1715](#): Enhance Ranger Hive Plugin to support authorization on Hive replication Tasks - repladmin permission
- [RANGER-1841](#): Audit log record for 'use dbName' hive command contains large number of tags
- [RANGER-1828](#): Good coding practice-add additional headers in ranger
- [RANGER-1833](#): Update Ranger to use 0.8.1 Atlas version
- [RANGER-1825](#): BUG-89494 : Ranger Tagsync start is failing
- [RANGER-1886](#): TagSync updates for changes in Atlas API
- [RANGER-1886](#): TagSync updates for changes in Atlas API - Part 2
- [RANGER-1883](#): Part 2 - All tag sources and sinks should use the same kerberos identity and reuse it
- [RANGER-1897](#): TagSync should replace use of V1 Atlas APIs with V2 APIs for efficient tag-download from Atlas
- [RANGER-1895](#): Simplify Storm dependencies
- [RANGER-1797](#): Tomcat Security Vulnerability Alert. The version of the tomcat for ranger should upgrade to 7.0.82.
- [RANGER-1906](#): Simplify Atlas plugin dependency management
- [RANGER-1898](#): Simplify Knox plugin dependency management
- [RANGER-1893](#): Update HBase dependency to 1.2.6
- [RANGER-1922](#): Simplify HBase dependencies
- [RANGER-1926](#): Upgrade Storm dependency to 1.1.1
- [RANGER-1905](#): Fix assembly creation - it's a black magic. As ranger-plugins-audit doesn't exclude guava from hadoop common, an older guava is leaked into embeddedwebserver classpath, which somehow filtered out by the assembly plugin Fix the atlas plugin assembly, so the guava is actually coming from the 'atlas-plugin's dependency, and not from the ranger-plugins-audit's dependency. And revert "RANGER-1858:ranger-admin start fail when policymgr\_http\_enabled equal to false."
- [RANGER-1984](#): Hbase audit log records may not show all tags associated with accessed column
- [RANGER-1999](#): Ranger policy engine updates to support list-of-values in access resource

- [RANGER-2000](#): Policy effective dates to support time-bound and temporary authorization
- [RANGER-1987](#): Upgrade to Storm 1.2.0 in Ranger
- [RANGER-2019](#): Handle upgrade scenario to rename the old ATLAS service def and use the new service def
- [RANGER-2019](#): Handle upgrade scenario to rename the old ATLAS service def and use the new service def
- [RANGER-2027](#): Evaluate grantor's group membership in the plugin for grant/revoke request
- [RANGER-2026](#): Update Hbase plugin to handle default namespace
- [RANGER-2018](#): Upgrade to Spring 4
- [RANGER-1994](#): Tomcat Security Vulnerability Alert. The version of the tomcat for ranger should upgrade to 7.0.85.
- [RANGER-2035](#): fix handling of null implClass in servicedefs (Oracle returns null for stored empty strings)
- [RANGER-2043](#): Ranger KMS KeyProvider and HSM KeyProvider should have more debug logs
- [RANGER-2115](#): Atlas service def creation failing in Ranger manual installation

## Spark

This release provides Spark 2.3.1 and the same list of patches listed for HDP 3.0.0.

- [SPARK-24495](#): SortMergeJoin with duplicate keys wrong results.
- [SPARK-207](#): Remove hardcoded FS scheme from Spark archive.

(Backport from 2.3.2)

- [SPARK-24455](#): fix typo in TaskSchedulerImpl comment.
- [SPARK-24369](#): Correct handling for multiple distinct aggregations having the same argument set.
- [SPARK-24468](#): Handle negative scale when adjusting precision for decimal operations.
- [SPARK-23732](#): Fix source links in generated scaladoc.
- [SPARK-24502](#): flaky test: UnsafeRowSerializerSuite.
- [SPARK-24531](#): Remove version 2.2.0 from testing versions in HiveExternalCatalogVersionsSuite.
- [SPARK-24506](#): Add UI filters to tabs added after binding.
- [SPARK-23754](#): Move UDF stop iteration wrapping from driver to executor.
- Remove unnecessary sort in UnsafeInMemorySorterSuite.
- Fix typo in serializer exception.
- Revert [SPARK-21743](#) top-most limit should not cause memory leak.
- [SPARK-24531](#): Replace 2.3.0 version with 2.3.1.

(Backport from 2.4)

- [SPARK-21783](#): Turn on `native` ORC impl and PPD by default.
- [SPARK-23456](#): Turn on `native` ORC impl and PPD by default.
- [SPARK-23228](#): Add Python Created jsparkSession to JVM's defaultSession.
- [SPARK-23510](#): Support Hive 2.2 and Hive 2.3 metastore.
- [SPARK-23518](#): Avoid metastore access when the users only want to read and write data frames.
- [SPARK-23635](#): Spark executor env variable is overwritten by same name AM env variable.
- [SPARK-23787](#): Fix file download test in SparkSubmitSuite for Hadoop 2.9..
- [SPARK-23355](#) convertMetastore should not ignore table properties.
- [SPARK-24110](#) Avoid UGI.loginUserFromKeytab in STS.
- [SPARK-24149](#): Retrieve all federated namespaces tokens.
- [SPARK-24209](#): Automatic retrieve proxyBase from Knox headers.
- [SPARK-24312](#): Upgrade to 2.3.3 for Hive Metastore Client 2.3.
- [SPARK-24377](#): make --py-files work in non pyspark application.
- [SPARK-24479](#): Added config for registering streamingQueryListeners.
- [SPARK-24518](#): Using Hadoop credential provider API to store password.
- [SPARK-24660](#): Show correct error pages when downloading logs in SHS.

## Livy

This release provides Livy 0.5.0 and no additional patches.

HDP 3.0.0 provided Livy 0.5.0 with no additional Apache patches.

## Sqoop

This release provides Sqoop 1.4.7 with no additional Apache patches.

HDP 3.0.0 provided Sqoop 1.4.7 with no additional Apache patches.

## Storm

This release provides Storm 1.2.1 and the same list of patches listed for HDP 3.0.0.

HDP 3.0.0 provided Storm 1.2.1 and the following Apache patches:

- [STORM-1839](#): Kinesis Spout
- [STORM-1694](#): Kafka Spout Trident Implementation Using New Kafka Consumer API
- [STORM-1960](#): Add CORS support to STORM UI Rest api
- [STORM-1959](#): Add missing license header to KafkaPartitionOffsetLag
- [STORM-1950](#): Change response json of "Topology Lag" REST API to keyed by spoutId, topic, partition.
- [STORM-1833](#): Simple equi-join in storm-sql standalone mode
- [STORM-1866](#): Update Resource Aware Scheduler Documentation
- [STORM-1930](#): Kafka New Client API - Support for Topic Wildcards
- [STORM-1924](#): Adding conf options for Persistent Word Count Topology
- [STORM-1956](#): Disabling Backpressure by default
- [STORM-1934](#): Fix race condition between sync-supervisor and sync-processes
- [STORM-1919](#): Introduce FilterBolt on storm-redis
- [STORM-1742](#): More accurate 'complete latency'
- [STORM-1945](#): Fix NPE bugs on topology spout lag for storm-kafka-monitor
- [STORM-1888](#): add description for shell command
- [STORM-1902](#): add a simple & flexible FileNameFormat for storm-hdfs
- [STORM-1914](#): Storm Kafka Field Topic Selector
- [STORM-1907](#): PartitionedTridentSpoutExecutor has incompatible types that cause ClassCastException
- [STORM-1925](#): Remove Nimbus thrift call from Nimbus itself
- [STORM-1909](#): Update HDFS spout documentation
- [STORM-1136](#): Command line module to return kafka spout offsets lag and display in storm ui
- [STORM-1911](#): IClusterMetricsConsumer should use seconds to timestamp unit
- [STORM-1849](#): HDFSFileTopology should use the 3rd argument as topologyName
- [STORM-1906](#): Window count/length of zero should be disallowed
- [STORM-1893](#): Support OpenTSDB for storing timeseries data.
- [STORM-1723](#): Introduce ClusterMetricsConsumer
- [STORM-1700](#): Introduce 'whitelist' / 'blacklist' option to MetricsConsumer
- [STORM-1698](#): Asynchronous MetricsConsumerBolt
- [STORM-1705](#): Cap number of retries for a failed message
- [STORM-1884](#): Prioritize pendingPrepare over pendingCommit
- [STORM-1575](#): fix TwitterSampleSpout NPE on close
- [STORM-1874](#): Update logger private permissions
- [STORM-1865](#): update command line client document
- [STORM-1771](#): HiveState should flushAndClose before closing old or idle Hive connections
- [STORM-1882](#): Expose TextFileReader public
- [STORM-1873](#): Implement alternative behaviour for late tuples
- [STORM-1719](#): Introduce REST API: Topology metric stats for stream

- [STORM-1887](#): Fixed help message for set\_log\_level command
- [STORM-1878](#): Flux can now handle IStatefulBolts
- [STORM-1864](#): StormSubmitter should throw respective exceptions and log respective errors for registered submitter hook invocation
- [STORM-1868](#): Modify TridentKafkaWordCount to run in distributed mode
- [STORM-1859](#): Ack late tuples in windowed mode
- [STORM-1851](#): Fix default nimbus impersonation authorizer config
- [STORM-1848](#): Make KafkaMessageId and Partition serializable to support
- [STORM-1862](#): Flux ShellSpout and ShellBolt can't emit to named streams
- [STORM-1728](#): TransactionalTridentKafkaSpout error
- [STORM-1850](#): State Checkpointing Documentation update
- [STORM-1674](#): Idle KafkaSpout consumes more bandwidth than needed
- [STORM-1842](#): Forward references in storm.thrift cause tooling issues
- [STORM-1730](#): LocalCluster#shutdown() does not terminate all storm threads/thread pools.
- [STORM-1841](#): Address a few minor issues in windowing and doc
- [STORM-1709](#): Added group by support in storm sql standalone mode
- [STORM-1720](#): Support GEO in storm-redis

## Tez

This release provides Tez 0.9.1 with no additional Apache patches.

HDP 3.0.0 provided Tez 0.9.1 with no additional Apache patches.

## Zeppelin

This release provides Zeppelin 0.8.0 with no additional Apache patches.

HDP 3.0.0 provided Zeppelin 0.8.0 with no additional Apache patches.

## ZooKeeper

This release provides ZooKeeper 3.4.6 with the following Apache patches.

- [ZOOKEEPER-2662](#): Export a metric for txn log sync times
- [ZOOKEEPER-3019](#): Add a metric to track number of slow fsyncs

HDP 3.0.0 provided ZooKeeper 3.4.6 and the following Apache patches:

- [ZOOKEEPER-2141](#): ACL cache in DataTree never removes entries
- [ZOOKEEPER-1897](#): ZK Shell/Cli not processing commands (stack via michim)
- [ZOOKEEPER-1926](#): Unit tests should only use build/test/data for data
- [ZOOKEEPER-2064](#): Prevent resource leak in various classes
- [ZOOKEEPER-2053](#): Make scripts use ZOOKEEPER\_HOME if set.
- [ZOOKEEPER-1876](#): Add support for installing windows services in .cmd scripts
- [ZOOKEEPER-1848](#): [WINDOWS] Java NIO socket channels does not work with Windows ipv6 on JDK6
- [ZOOKEEPER-1574](#): Set svn:eol-style property to native.
- [ZOOKEEPER-1575](#): adding .gitattributes to prevent CRLF and LF mismatches for source and text files (Raja Aluri via michim)
- [ZOOKEEPER-1888](#): ZkCli.cmd commands fail with "'java' is not recognized as an internal or external command" (Ivan Mitic via michim)
- [ZOOKEEPER-1791](#): ZooKeeper package includes unnecessary jars that are part of the package. (mahadev)
- [ZOOKEEPER-1506](#): Re-try DNS hostname -> IP resolution if node connection fails
- [ZOOKEEPER-1952](#): Default log directory and file name can be changed
- [ZOOKEEPER-706](#): large numbers of watches can cause session re-establishment to fail (Chris Thunes)
- [ZOOKEEPER-1576](#): Zookeeper cluster - failed to connect to cluster if one of the provided IPs causes java.net.UnknownHostException (Edward Ribeiro via camille)



## Fixed Common Vulnerabilities and Exposures

This section covers all Common Vulnerabilities and Exposures (CVE) that are addressed in this release.

### CVE-2018-8008

Component: Storm

Summary: Apache Storm arbitrary file write vulnerability.

Severity: Important

Vendor: The Apache Software Foundation

Versions Affected: HDP 3.0.0, HDP 2.6.5 and earlier

Impact: Apache Storm version 1.0.6 and earlier, 1.2.1 and earlier, and version 1.1.2 and earlier expose an arbitrary file write vulnerability, that can be achieved using a specially crafted zip archive (affects other archives as well, bzip2, tar, xz, war, cpio, 7z), that holds path traversal filenames. So when the filename gets concatenated to the target extraction directory, the final path ends up outside of the target folder.

Recommended Action: Upgrade to HDP 3.0.1.

### CVE-2018-8009

Summary: Zip Slip Vulnerability - Apache Hadoop distributed cache archive vulnerability

Severity: Critical

Vendor: Hortonworks

Versions Affected: HDP 3.0.0

Users Affected: Users who run Mapreduce jobs.

Impact: Zip Slip is a widespread arbitrary file overwrite critical vulnerability, which typically results in remote command execution. It was discovered and responsibly disclosed by the Snyk Security team.

Recommended Action: Upgrade to HDP 3.0.1+.

### CVE-2018-12536

Summary: InvalidPathException message

Severity: Moderate

Vendor: Hortonworks

Versions Affected: HDP 3.0.0

Users Affected: Users who use the Spark user interfaces.

Impact: When an intentionally bad query arrives, the message included in the error response can reveal the full server path to the requesting system.

Recommended Action: Upgrade to HDP 3.0.1.

### CVE-2018-11778

Summary: Apache Ranger Stack based buffer overflow

Severity: Critical

Vendor: Hortonworks

Versions Affected: HDP 2.3/2.4/2.5/2.6/3.0 versions including Apache Ranger versions 0.5.x/0.6.x/0.7.x/1.1.0

Users Affected: Environments that use unix authentication service.

Impact: Ranger UnixAuthenticationService should properly handle user input to avoid Stack-based buffer overflow.

Recommended Action: Upgrade to HDP 3.0.1+.

## Fixed Issues

Fixed issues represents selected issues that were previously logged via Hortonworks Support, but are now addressed in the current release. These issues may have been reported in previous versions within the Known Issues section; meaning they were reported by customers or identified by Hortonworks Quality Engineering team.

### Incorrect Results

Hortonworks Bug ID	Apache JIRA	Summary
BUG-108283	<a href="#">RANGER-2169</a>	Create unique index on service and name column of x_policy table
BUG-108346	<a href="#">KNOX-1398</a>	WebHDFS and WebHBase APIs listed incorrectly as services that need whitelist enforcement
BUG-108348	<a href="#">KNOX-1392</a> , <a href="#">KNOX-1393</a> , <a href="#">KNOX-1395</a>	Remove use of InetAddress in Default Whitelist Creation
BUG-108351	<a href="#">KNOX-1386</a>	Fix NiFi service discovery
BUG-108352	<a href="#">KNOX-1385</a>	Whitelist test inconsistencies
BUG-108802	<a href="#">KNOX-1412</a>	Knox fails while executing a `curl` command using Content-Type: text/xml" "Content-Encoding: gzip"
BUG-108882	<a href="#">RANGER-2187</a>	External Group search fails on Ranger UI when installed with postgres
BUG-109034	<a href="#">RANGER-2114</a>	Backport RANGER-2114 on Atlantic-M05
BUG-109385	<a href="#">HIVE-20406</a>	Nested Coalesce giving incorrect results
BUG-109887	<a href="#">HIVE-20187</a>	Backport HIVE-20187
BUG-95462	<a href="#">HDFS-13205</a> , <a href="#">RANGER-2183</a>	Incorrect path is passed to checkPermission during authorization of file under a snapshot (specifically under a subdir) after original subdir is deleted

### Other

Hortonworks Bug ID	Apache JIRA	Summary
BUG-102276	<a href="#">HIVE-19739</a>	Use the checkpoint on table / partition to create relevant repl load tasks.
BUG-102279	<a href="#">HIVE-19739</a>	Repl load should not fail if Target Db already exists.
BUG-103052	<a href="#">HIVE-19564</a>	Vectorization: Backport HIVE-19564 to Atlantic-M05 --> Fix NULL / Wrong Results issues in Arithmetic Functions
BUG-105695	<a href="#">HIVE-20090</a>	TPCDS query64 has regressed in Atlantic by 5x times.
BUG-105963	<a href="#">HIVE-19408</a>	Show tables statement includes views and materialized views
BUG-106249	<a href="#">HIVE-20011</a>	Move away from append mode in proto logging hook
BUG-106294	<a href="#">HIVE-20028</a>	hive streaming fails with unable to abort invalid txns
BUG-106359	<a href="#">HIVE-4367</a>	enhance TRUNCATE syntax to drop data of external table
BUG-107015	N/A	TestReplicationScenariosAcidTables & TestReplicationOnHDFSEncryptedZones failing
BUG-107018	<a href="#">HIVE-19104</a>	TestHS2ImpersonationWithRemoteMS.testImpersonation failing
BUG-107020	<a href="#">HIVE-20299</a>	TestLlapSignerImpl.testSigning failing
BUG-107023	N/A	TestDynamicPartitionPruner.testSingleSourceMultipleFiltersOrdering1 failing with time out

Hortonworks Bug ID	Apache JIRA	Summary
BUG-107024	N/A	TestCustomPartitionVertex.testGetBytePayload failing due to timeout
BUG-107025	<a href="#">HIVE-19800</a>	TestUpgradeTool.testPostUpgrade failing
BUG-107026	N/A	TestWarehouseExternalDir.testManagedPaths failing
BUG-107062	N/A	TestSchemaToolForMetastore failing
BUG-107135	<a href="#">HIVE-20353</a>	Follow redirects when hive connects to a passive druid overlord/coordinator
BUG-107180	N/A	CLONE - UT error in storm: org.apache.storm.sql.hdfs.TestHdfsDataSourcesProvider
BUG-107264	N/A	TestMiniDruidCliDriver failing
BUG-107300	<a href="#">KNOX-1386</a>	Autodiscovery not using the correct NiFi properties
BUG-107398	<a href="#">HIVE-20151</a>	External table: exception while storing stats
BUG-107516	<a href="#">HIVE-19800</a>	PreUpgrade Tool: Support upgrade from >= 2.6.0
BUG-107547	<a href="#">HIVE-20135</a>	Backport HIVE-20135: Fix incompatible change in TimestampColumnVector to default to UTC
BUG-107548	<a href="#">HIVE-20090</a>	Backport HIVE-20090: Extend creation of semijoin reduction filters to be able to discover new opportunities
BUG-107553	N/A	Vectorization: Backport HIVE-19992 to Atlantic-M05: Follow-on to HIVE-19951 --> add call to SchemaEvolution.isOnlyImplicitConversion to disable encoded LLAP I/O for ORC only when data type conversion is not implicit
BUG-107572	N/A	Vectorization: Backport HIVE-20174 to Atlantic-M05: Fix NULL / Wrong Results issues in GROUP BY Aggregation Functions
BUG-107694	<a href="#">HIVE-20098</a>	Backport HIVE-20098: Statistics: NPE when getting Date column partition statistics
BUG-107695	<a href="#">HIVE-20152</a>	Backport HIVE-20152: reset db state, when repl dump fails, so rename table can be done
BUG-107713	N/A	Vectorization: Backport HIVE-20197 to Atlantic-M05: Add DECIMAL_64 testing, add Date/Interval/Timestamp arithmetic, and add more GROUP BY Aggregation tests
BUG-107768	<a href="#">HIVE-19360</a>	backport: CBO: Add an "optimizedSQL" to QueryPlan object
BUG-107866	<a href="#">ATLAS-2786</a>	Atlas HA fails to use all the zookeeper in the environment
BUG-108021	<a href="#">HIVE-20383</a>	provide the correct queue name to proto logger
BUG-108080	<a href="#">HIVE-20172</a>	Backport HIVE-20172: StatsUpdater failed with GSS Exception while trying to connect to remote metastore
BUG-108109	N/A	Disable flaky spark Tests
BUG-108201	<a href="#">HIVE-20207</a>	Vectorization: Backport HIVE-20207 to Atlantic-M05: Fix NULL / Wrong Results issues in Filter / Compare
BUG-108216	<a href="#">HIVE-20192</a>	HS2 with embedded metastore is leaking JDOPersistenceManager objects.
BUG-108263	N/A	Fix TestJdbcWithDBTokenStore.testTokenAuth, TestJdbcWithDBTokenStoreNoDoAs.testTokenAuth failures
BUG-108265	<a href="#">HIVE-20212</a>	Backport HIVE-20212: Hiveserver2 in http mode emitting metric default.General.open_connections incorrectly
BUG-108287	<a href="#">HIVE-20383</a>	Missing pre-hook and duplicate post hook events from hive proto events hook.
BUG-108345	<a href="#">HIVE-20245</a>	Vectorization: Backport HIVE-20245 to Atlantic-M05: Fix NULL / Wrong Results issues in BETWEEN / IN
BUG-108357	<a href="#">HIVE-20082</a> , <a href="#">HIVE-20204</a>	Decimal truncation for trailing zeros - port to 3.0-maint
BUG-108424	<a href="#">HIVE-20033</a>	PowerBI unable to connect to LLAP cluster with large Databases/Tables
BUG-108439	N/A	TestJdbcWithMiniLlapRow is failing on ptest
BUG-108440	N/A	TestHiveDruidQueryBasedInputFormat is failing on ptest

Hortonworks Bug ID	Apache JIRA	Summary
BUG-108444	<a href="#">HIVE-18201</a>	Subquery joins result in cross product instead of broadcast
BUG-108511	<a href="#">HIVE-20263</a>	Backport HIVE-20263: Typo in HiveReduceExpressionsWithStatsRule variable
BUG-108512	<a href="#">HIVE-20241</a>	Backport HIVE-20241: Support partitioning spec in CTAS statements
BUG-108513	<a href="#">HIVE-14493</a>	Backport HIVE-14493: Partitioning support for materialized views
BUG-108558	<a href="#">HIVE-20281</a>	Backport HIVE-20281: SharedWorkOptimizer fails with 'operator cache contents and actual plan differ'
BUG-108566	N/A	TPCDS query23 runs longer on hdp3.0 with varying runtimes.
BUG-108573	<a href="#">KNOX-1415</a>	Upgrade Jetty to 9.2.25
BUG-108586	N/A	Vectorization: Backport HIVE-20294 to Atlantic-M05: Fix NULL / Wrong Results issues in COALESCE / ELT
BUG-108642	N/A	Update golden files for TezDriver
BUG-108643	<a href="#">HIVE-20302</a>	Backport HIVE-20302: LLAP: non-vectorized execution in IO ignores virtual columns, including ROW_ID
BUG-108644	<a href="#">HIVE-20314</a>	Backport HIVE-20314: Include partition pruning in materialized view rewriting
BUG-108645	<a href="#">HIVE-20301</a>	Backport HIVE-20301: Enable vectorization for materialized view rewriting tests
BUG-108719	<a href="#">HIVE-20322</a>	FlakyTest: TestMiniDruidCliDriver
BUG-108720	N/A	Add missing Apache Hadoop 3.1.1 commits to Atlantic M05 (YARN/MR)
BUG-108816	<a href="#">HIVE-20383</a>	Hive: ProtoLoggingHook NPE on queue name field
BUG-108894	<a href="#">HIVE-20336</a>	Backport HIVE-20336: Masking and filtering policies for materialized views
BUG-108910	N/A	Tez History API fails with ClassCastException
BUG-108917	N/A	Backport Druid perf enhancements
BUG-108999	<a href="#">HIVE-20347</a>	Backport HIVE-20347: hive.optimize.sort.dynamic.partition should work with partitioned CTAS and MV
BUG-109060	<a href="#">HIVE-20364</a>	Update default for hive.map.aggr.hash.min.reduction
BUG-109168	N/A	TestHiveCli timing out (need HIVE-20149 backport)
BUG-109173	N/A	Backport CALCITE-2465: Enable use of materialized views for any planner
BUG-109245	<a href="#">HIVE-20379</a>	Backport HIVE-20379: Rewriting with partitioned materialized views may reference wrong column
BUG-109247	<a href="#">HIVE-20382</a>	Backport HIVE-20382: Materialized views: Introduce heuristic to favour incremental rebuild
BUG-109272	<a href="#">HIVE-20345</a>	Backport HIVE-20345: Drop database may hang if the tables get deleted from a different call
BUG-109311	<a href="#">HIVE-20391</a>	Backport HIVE-20391: HiveAggregateReduceFunctionsRule may infer wrong return type when decomposing aggregate function
BUG-109313	N/A	Backport HIVE-17979 : Tez: Improve ReduceRecordSource passDownKey copying
BUG-109319	N/A	Backport HIVE-20177: Vectorization: Reduce KeyWrapper allocation in GroupBy Streaming mode
BUG-109321	N/A	Backport HIVE-20153: Count and Sum UDF consume more memory in Hive 2+
BUG-109703	<a href="#">RANGER-2167</a>	Upgrade to Apache parent pom version 20
BUG-110109	N/A	Remove unicode from superset
BUG-110124	<a href="#">YARN-8726</a>	YARN UI2 is not accessible
BUG-110428	<a href="#">HIVE-20450</a>	Add replication test for LOAD command on ACID table.
BUG-89982	<a href="#">HIVE-20357</a>	Forward port azure db scripts to atlantic

Hortonworks Bug ID	Apache JIRA	Summary
BUG-92690	<a href="#">HIVE-19267</a>	Generate new event for ACID Write operations
BUG-92700	<a href="#">HIVE-19340</a>	Disable timeout of transactions opened by replication task at target cluster
BUG-95678	<a href="#">HIVE-17683</a>	Add explain locks <sql> command
BUG-98176	<a href="#">HIVE-20337</a>	TPCDS query16 compilation fails on cn105 when CachedStore is used. ObjectStore works fine.
BUG-99461	<a href="#">HIVE-20337</a>	Queries returning incorrect results using CachedStore. ObjectStore works fine.
BUG-110970	N/A	For long-running SparkSQL jobs on a Kerberized cluster, some JDBC clients may randomly fail with a "no token found in cache" error after the delegation token expiry period.

## Performance

Hortonworks Bug ID	Apache JIRA	Summary
BUG-105965	N/A	Adding constraints on TPCDS schema, changes query plan for query93 and degrades its perf by 4x times
BUG-107255	<a href="#">HIVE-20209</a>	Metastore connection fail 1st attempt in repl dump
BUG-107475	<a href="#">RANGER-2162</a>	Rangeradmin stuck with db connection issue.
BUG-107932	<a href="#">RANGER-2158</a>	[Atlantic-M05] - RANGER-2158 Performance improvement for REST API calls
BUG-107990	<a href="#">HIVE-20300</a>	VectorFileSinkArrowOperator
BUG-108108	<a href="#">HADOOP-15598</a>	DataChecksum calculate checksum is contented on hashtable synchronization
BUG-108281	<a href="#">RANGER-2165</a>	Address JPA Cache issue when policies Create, Update and Delete are done via REST API in Apache Ranger admin
BUG-108282	<a href="#">RANGER-2168</a>	Add service admin user through service config
BUG-108401	<a href="#">RANGER-2173</a>	RANGER-2173: Optimize Trie construction and Policy lookup
BUG-108422	<a href="#">HADOOP-15547</a>	WASB: improve listStatus performance
BUG-108515	<a href="#">HIVE-20290</a>	Lazy initialize ArrowColumnarBatchSerDe so it doesn't allocate buffers during GetSplits
BUG-108638	<a href="#">HIVE-20316</a>	Skip external table file listing for create table event
BUG-108742	<a href="#">HIVE-18624</a>	Hive parser is very slow in parsing nested expressions
BUG-109693	<a href="#">RANGER-2188</a>	Support multiple threads to build Trie and on-lookup post-setup for Trie nodes
BUG-109694	<a href="#">RANGER-2191</a>	Update ranger-tool with new options to control Trie
BUG-109702	<a href="#">RANGER-2186</a>	Increment service-specific policy and tag versions after...
BUG-109720	<a href="#">RANGER-2196</a>	Ensure that any explicit threads used by Ranger are marked as daemon threads

## Potential Data Loss

Hortonworks Bug ID	Apache JIRA	Summary
BUG-104681	<a href="#">HIVE-19891</a>	insert from select overwrites the partition location path

## Security

Hortonworks Bug ID	Apache JIRA	Summary
BUG-106840	N/A	CVE-2018-8009: Zip Slip Vulnerability - Apache Hadoop distributed cache archive vulnerability
BUG-107549	N/A	druid bolt fails with NoSuchMethodError

Hortonworks Bug ID	Apache JIRA	Summary
BUG-107610	<a href="#">ATLAS-2788</a>	Add policy permissions for tagsync user
BUG-107863	<a href="#">STORM-3052</a>	Backport STORM-3052: Let blobs un archive (HDP 3.0-maint)
BUG-108451	<a href="#">KNOX-1416</a>	Upgrade Spring Framework to 4.3.17
BUG-108574	N/A	Upgrade Bouncy Castle Dependency from Pac4J
BUG-108749	<a href="#">ZEPPELIN-3686</a>	Update org.eclipse.jetty version to 9.4.11.v20180605 to avoid CVE issues
BUG-108874	N/A	Upgrade of Jackson Databind to 2.9.5 in Oozie
BUG-108997	N/A	BlackDuck for HWC
BUG-109155	N/A	Address Storm black-duck scan issues
BUG-109209	<a href="#">ATLAS-2821</a>	Update old atlas-simple-authz-policy.json (m0) file with relationship permission attributes

### Stability

Hortonworks Bug ID	Apache JIRA	Summary
BUG-101700	<a href="#">SPARK-23007</a>	Add read schema suite for file-based data sources
BUG-102162	<a href="#">YARN-8595</a>	[UI2] Container diagnostic missing from container page
BUG-102349	<a href="#">YARN-8593</a> , <a href="#">YARN-8594</a>	[UI2] Show the current logged in user in UI2
BUG-102389	<a href="#">YARN-8620</a>	[UI2] YARN Services UI new submission failures are not debuggable
BUG-102493	<a href="#">YARN-8608</a>	[UI2] No information available per application appAttempt about 'Total Outstanding Resource Requests'
BUG-103056	N/A	[YARN-8301] Add documentation for yarn service upgrade
BUG-104220	<a href="#">RANGER-2150</a>	Unit test coverage for org.apache.ranger.biz.XUserMgr and UserMgr
BUG-104727	N/A	[YARN-8403] NM log should print Exception at ERROR log level
BUG-104844	<a href="#">YARN-8407</a>	[YARN-8407] Container launch exception in AM log should be printed in ERROR level
BUG-105493	<a href="#">YARN-8429</a>	[YARN-8429] Improve diagnostic message when artifact is not set properly
BUG-106203	<a href="#">YARN-8603</a>	[UI2] Latest run application should be listed first in the RM UI
BUG-106504	N/A	Update `spark2:atlantic` with Spark 2.3.2 RC5
BUG-106619	N/A	Zengine UT failure: IllegalStateException: Queue configuration missing child queue names for root
BUG-107236	<a href="#">ATLAS-2793</a> , <a href="#">ATLAS-2794</a> , <a href="#">RANGER-2166</a>	Tagsync start fails with CNF error for org/codehaus/jackson/xc/JaxbAnnotationIntrospector when using Atlas-REST as source
BUG-107241	<a href="#">HIVE-19860</a>	HiveServer2 ObjectInspectorFactory memory leak with cachedUnionStructObjectInspector
BUG-107347	N/A	Zeppelin, set right URL for jdbc for different transport modes
BUG-107413	<a href="#">SPARK-23679</a>	Spark Web UI in master=yarn mode does not proxy/redirect with YARN HA (secured)
BUG-107477	<a href="#">YARN-8418</a>	[YARN-8418] App local logs could leaked if log aggregation fails to initialize for the app
BUG-107576	<a href="#">HIVE-20069</a>	Backport HIVE-20069: Fix reoptimization in case of DPP and Semijoin optimization
BUG-107632	N/A	[YARN-8545] YARN native service should return container if launch failed
BUG-107700	<a href="#">HIVE-20203</a>	Arrow SerDe leaks a DirectByteBuffer
BUG-107725	N/A	Hive warehouse directory creation issue at blueprint cluster creation

Hortonworks Bug ID	Apache JIRA	Summary
BUG-107871	N/A	[Service Discovery] Add Support for discovering beacon service
BUG-107922	<a href="#">HADOOP-15593</a>	Backport HADOOP-15593 to Atlantic-M05
BUG-107945	<a href="#">RANGER-2160</a>	'Email Address' search is not working in user listing page,userRoles filters also needs to be improved.
BUG-107955	<a href="#">YARN-8571</a>	op/cli commands to a running service is failing in a secure cluster
BUG-107965	<a href="#">RANGER-2164</a>	Ranger to add default atlas policy for rangertagsync user.
BUG-108112	N/A	backport HIVE-19935 (race in WM resulting in a session dying)
BUG-108146	<a href="#">RANGER-2182</a>	Handle upgrade scenario since atlas-service def is added with new resources for relationship
BUG-108209	N/A	[ATSV2] [YARN-8591] NPE error while checking for entity acl in non-secure cluster
BUG-108210	<a href="#">YARN-8592</a>	[UI2] Resource Manager Web UI V2 shows a blank page in windows OS And Chrome browser
BUG-108283	<a href="#">RANGER-2169</a>	Create unique index on service and name column of x_policy table
BUG-108349	<a href="#">KNOX-1391</a>	Fix IllegalStateException when accessing Admin UI
BUG-108350	<a href="#">KNOX-1387</a> , <a href="#">KNOX-1390</a>	Fix NPE in AclsAuthorizationFilter
BUG-108354	<a href="#">HIVE-20244</a> , <a href="#">HIVE-20247</a> , <a href="#">HIVE-20249</a>	backport LLAP error handling changes
BUG-108407	<a href="#">RANGER-2171</a>	RANGER-2171: Unit Test cases to cover policy operations from service admin user
BUG-108583	N/A	Issue with tutorial notebooks
BUG-108624	N/A	ATSV2 HBase Application does not appear when HBase app is stopped [System Service Mode]
BUG-108632	<a href="#">HIVE-20312</a>	HWC: Assign Arrow ChildAllocator per-task to fail-fast on memory leak
BUG-108762	<a href="#">SPARK-24948</a>	Spark History Server filters some applications on ADL
BUG-108833	N/A	Upgrade bouncy castle version to 1.6.0 to fix CVE issue
BUG-108915	N/A	Upgrade Jetty version to fix CVE issue in Spark
BUG-109034	<a href="#">RANGER-2114</a>	Backport RANGER-2114 on Atlantic-M05
BUG-109183	N/A	[YARN-8667] Container Relaunch fails with "find: File system loop detected;"
BUG-109184	N/A	Add yarn.service.container-failure.retry.max for ats-hbase app.
BUG-109195	<a href="#">SPARK-25051</a>	Correctness issue with outer joins for Spark 2.3
BUG-109198	N/A	Provide work around handle BUG-108844 impact on YARN
BUG-109312	<a href="#">HIVE-20393</a>	Semijoin Reduction : markSemiJoinForDPP behaves inconsistently
BUG-109384	N/A	[YARN-8679] Gridmix of HDP-3.0 is stuck
BUG-109581	<a href="#">ZEPPELIN-3741</a>	Do not clear "Authorization" header if Z-server is running behind proxy
BUG-109697	<a href="#">RANGER-2180</a>	Handle token replacement correctly when token is not defined in the request context
BUG-109720	<a href="#">RANGER-2196</a>	Ensure that any explicit threads used by Ranger are marked as daemon threads
BUG-109777	<a href="#">RANGER-2197</a>	Delegate Admin is not able to create policy

### Supportability

Hortonworks Bug ID	Apache JIRA	Summary
BUG-107437	<a href="#">ZOOKEEPER-2662</a> , <a href="#">ZOOKEEPER-3019</a>	Port JMX metrics for Zookeeper I/O to Atlantic-Next
BUG-109698	<a href="#">RANGER-2152</a>	Incorrect debugging information in RangerPluginClassLoader.java

Hortonworks Bug ID	Apache JIRA	Summary
BUG-109700	<a href="#">RANGER-2177</a>	Handle validations for duplicate configuration item during service create/edit
BUG-109701	<a href="#">RANGER-2153</a>	Supply the function of reverting policy history version.
BUG-30983	<a href="#">HIVE-14898</a>	HiveServer2 error message from Knox in logs - benign error confuses users

## Upgrade

Hortonworks Bug ID	Apache JIRA	Summary
BUG-108146	<a href="#">RANGER-2182</a>	Handle upgrade scenario since atlas-service def is added with new resources for relationship
BUG-109209	<a href="#">ATLAS-2821</a>	Update old atlas-simple-authz-policy.json (m0) file with relationship permission attributes
BUG-109696	<a href="#">RANGER-2184</a>	Update RangerAtlas authorization to authorize add/update/remove of relationships
BUG-109728	N/A	phoenix-sqlline launch fails on upgraded cluster

## Usability

Hortonworks Bug ID	Apache JIRA	Summary
BUG-103643	<a href="#">RANGER-2161</a>	it will be good to mark exclude from deny and allow condition with some background color
BUG-104455	<a href="#">HIVE-19694</a>	create materialized view should validate if view already exists
BUG-106469	<a href="#">HIVE-20044</a>	[SPARK_LLAP] fixed char type output via spark + HDWC is inconsistent with that via HS2 beeline
BUG-106681	<a href="#">OOZIE-3328</a>	Reenable hive action in Oozie
BUG-106688	<a href="#">HIVE-20341</a>	Druid Needs Explicit CASTs from Timestamp to STRING when the output of timestamp function is used as String
BUG-106818	<a href="#">ATLAS-2785</a>	import_hive.sh with "-t" option imports all metadata instead of just that table.
BUG-106266	<a href="#">OOZIE-2769</a> <a href="#">OOZIE-3085</a> <a href="#">OOZIE-3156</a> <a href="#">OOZIE-3183</a>	
BUG-107236	<a href="#">ATLAS-2793</a> , <a href="#">ATLAS-2794</a> , <a href="#">RANGER-2166</a>	Tagsync start fails with CNF error for org/codehaus/jackson/xc/JaxbAnnotationIntrospector when using Atlas-REST as source
BUG-107241	<a href="#">HIVE-19860</a>	HiveServer2 ObjectInspectorFactory memory leak with cachedUnionStructObjectInspector
BUG-107742	NA	Update some Ambari defaults based on 3.0-EA
BUG-107945	<a href="#">RANGER-2160</a>	'Email Address' search is not working in user listing page, userRoles filters also needs to be improved.
BUG-107965	<a href="#">RANGER-2164</a>	Ranger to add default atlas policy for rangertagsync user.
BUG-108114	N/A	Backport HIVE-20240
BUG-108282	<a href="#">RANGER-2168</a>	Add service admin user through service config
BUG-108312	N/A	Oozie doesn't work out of the box with Spark2
BUG-108347	<a href="#">KNOX-1383</a> , <a href="#">KNOX-1396</a> , <a href="#">KNOX-1397</a>	Admin UI Fixes for Topology List Refresh and URL Validation
BUG-108630	<a href="#">HIVE-20277</a>	Vectorization: Backport HIVE-20277 to Atlantic-M05: Vectorization: Case expressions that return BOOLEAN are not supported for FILTER
BUG-108738	N/A	HWC: Python API is missing "executeUpdate"
BUG-108882	<a href="#">RANGER-2187</a>	External Group search fails on Ranger UI when installed with postgres
BUG-109034	<a href="#">RANGER-2114</a>	Backport RANGER-2114 on Atlantic-M05



Hortonworks Bug ID	Apache JIRA	Summary
BUG-109290	N/A	HWC: HiveWarehouseSession.table(<tableName>) doesn't accept unqualified table name
BUG-109309	N/A	HWC: HiveWarehouseSession.dropTable API is broken
BUG-109698	<a href="#">RANGER-2152</a>	Incorrect debugging information in RangerPluginClassLoader.java
BUG-109700	<a href="#">RANGER-2177</a>	Handle validations for duplicate configuration item during service create/edit
BUG-109701	<a href="#">RANGER-2153</a>	Supply the function of reverting policy history version.
BUG-109777	<a href="#">RANGER-2197</a>	Delegate Admin is not able to create policy
BUG-109784	N/A	Misleading documentation to enable SSL on Spark2 History Server
BUG-111169	N/A	Oozie hive spark jobs failing because spark still use heat type delegation tokens
BUG-95269	<a href="#">YARN-8318</a>	[UI2] Few Missing information on Component Instance Information page
BUG-95462	<a href="#">HDFS-13205</a> , <a href="#">RANGER-2183</a>	Incorrect path is passed to checkPermission during authorization of file under a snapshot (specifically under a subdir) after original subdir is deleted
BUG-95678	<a href="#">HIVE-17683</a>	Add explain locks <sql> command

## Known Issues

Summary of known issues for this release.

Hortonworks Bug ID	Apache JIRA	Apache component	Summary
BUG-79238	N/A	Documentation, HBase, HDFS, Hive, MapReduce, Zookeeper	<p>Description of the problem or behavior</p> <p>SSL is deprecated and its use in production is not recommended. Use TLS.</p> <p>Workaround</p> <p>In Ambari: Use <code>ssl.enabled.protocols=TLSv1 TLSv1.1 TLSv1.2</code> and <code>security.server.disabled.protocols=SSL SSLv2 SSLv3</code>. For help configuring TLS for other components, contact customer support. Documentation will be provided in a future release.</p>

Hortonworks Bug ID	Apache JIRA	Apache component	Summary
BUG-106494	N/A	Documentation, Hive	<p>Description of Problem</p> <p>When you partition a Hive column of type double, if the column value is 0.0, the actual partition directory is created as "0". An AIOB exception occurs.</p> <p>Associated error message</p> <pre> 2018-06-28T22:43:55,498 ERROR 441773a0-851c-4b25-9e47-729183946a26 main exec.StatsTask: Failed to run stats task org.apache.hadoop.hive.ql.metadata.HiveException: java.lang.IndexOutOfBoundsException: Index: 8, Size: 8 at org.apache.hadoop.hive.ql.metadata.Hive.setParti ~hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.stats.ColStatsProcesso ~hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.stats.ColStatsProcesso ~hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.exec.StatsTask.execut hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.exec.Task.executeTask hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.exec.TaskRunner.runSe hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.Driver.launchTask(Driv hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.Driver.execute(Driver hive-exec-4.0.0-SNAPSHOT.jar:4.0.0- SNAPSHOT at org.apache.hadoop.hive.ql.Driver.run </pre> <p>Workaround</p> <p>Do not partition columns of type double.</p>
BUG-106379	N/A	Documentation, Hive	<p>Description of the Problem</p> <p>The upgrade process fails to perform necessary compaction of ACID tables and can cause permanent data loss.</p> <p>Workaround</p> <p>If you have ACID tables in your Hive metastore, enable ACID operations in Ambari or set Hive configuration properties to enable ACID. If ACID operations are disabled, the upgrade process does not convert ACID tables. This causes permanent loss of data; you cannot recover data in your ACID tables later.</p>

Hortonworks Bug ID	Apache JIRA	Apache component	Summary
BUG-106286	N/A	Documentation, Hive	<p>Description of the Problem</p> <p>The upgrade process might fail to make a backup of the Hive metastore, which is critically important.</p> <p>Workaround</p> <p>Manually make a manual backup of your Hive metastore database before upgrading. Making a backup is especially important if you did not use Ambari to install Hive and create the metastore database, but highly recommended in all cases. Ambari might not have the necessary permissions to perform the backup automatically. The upgrade can succeed even if the backup fails, so having a backup is critically important.</p>
BUG-101082	N/A	Documentation, Hive	<p>Description of the problem or behavior</p> <p>When running Beeline in batch mode, queries killed by the Workload Management process can on rare occasions mistakenly return success on the command line.</p> <p>Workaround</p> <p>There is currently no workaround.</p>
BUG-103495	<a href="#">HBASE-20634</a> , <a href="#">HBASE-20680</a> , <a href="#">HBASE-20700</a>	HBase	<p>Description of the problem or behavior</p> <p>Because the region assignment is refactored in HBase, there are unclear issues that may affect the stability of this feature. If you rely on RegionServer Groups feature, you are recommended to wait until a future HDP 3.x release, which will return the stability of this features as it was available in HBase 1.x/HDP 2.x releases.</p> <p>Workaround</p> <p>There is currently no workaround.</p>
BUG-98727	N/A	HBase	<p>Description of the problem or behavior</p> <p>Because the region assignment is refactored in HBase, there are unclear issues that may affect the stability of this feature. If you rely on Region replication feature, you are recommended to wait until a future HDP 3.x release, which will return the stability of this features as it was available in HBase 1.x/HDP 2.x releases.</p> <p>Workaround</p> <p>There is currently no workaround.</p>
BUG-105983	N/A	HBase	<p>Description of the problem or behavior</p> <p>An HBase service (Master or RegionServer) stops participating with the rest of the HBase cluster.</p> <p>Associated error message</p> <p>The service's log contains stack traces that contain "Kerberos principal name does NOT have the expected hostname part..."</p> <p>Workaround</p> <p>Retrying the connection solves the problem.</p>

Hortonworks Bug ID	Apache JIRA	Apache component	Summary
BUG-94954	<a href="#">HBASE-20552</a>	HBase	<p>Description of the problem or behavior</p> <p>After a rolling restart of HBase, the HBase master may not correctly assign out all Regions to the cluster.</p> <p>Associated error message</p> <p>TTThere are regions in transition, including hbase:meta, which result in "Region is not online on RegionServer" messages on Master or RegionServer or messages around errors in ServerCrashProcedure in the Master.</p> <p>Workaround</p> <p>Restart the HBase Master.</p>
BUG-96402	<a href="#">HIVE-18687</a>	Hive	<p>Description of the problem or behavior</p> <p>When HiveServer2 is running in HA (high-availability) mode in HDP 3.0.0, resource plans are loaded in-memory by all HiveServer2 instances. If a client makes changes to a resource plan, the changes are reflected (pushed) only in the HiveServer2 to which the client is connected.</p> <p>Workaround</p> <p>In order for the resource plan changes to be reflected on all HiveServer2 instances, all HiveServer2 instances has to be restarted so that they can reload the resource plan from metastore.</p>
BUG-88614	N/A	Hive	<p>Description of the problem or behavior</p> <p>RDMBS schema for Hive metastore contains an index HL_TXNID_INDEX defined as</p> <pre>CREATE INDEX HL_TXNID_INDEX ON HIVE_LOCKS USING hash (HL_TXNID);</pre> <p>Hash indexes are not recommended by PostgreSQL. For more information, see <a href="https://www.postgresql.org/docs/9.4/static/indexes-types.html">https://www.postgresql.org/docs/9.4/static/indexes-types.html</a></p> <p>Workaround</p> <p>It's recommended that this index is changed to type BTREE.</p>
BUG-120655	N/A	Hive	<p>Description of the problem or behavior</p> <p>The Hive Warehouse Connector does not support non-ORC file formats for writes.</p> <p>Workaround</p> <p>There is currently no workaround.</p>
BUG-60904	<a href="#">KNOX-823</a>	Knox	<p>Description of the problem or behavior</p> <p>When Ambari is being proxied by Apache Knox, the QuickLinks are not rewritten to go back through the gateway. If all access to Ambari is through Knox in the deployment, the new Ambari QuickLink profile may be used to hide and/or change URLs to go through Knox permanently. Future release will make these reflect the gateway appropriately.</p> <p>Workaround</p> <p>There is currently no workaround.</p>

Hortonworks Bug ID	Apache JIRA	Apache component	Summary
BUG-107399	N/A	Knox	<p>Description of the problem or behavior</p> <p>After upgrade from previous HDP versions, certain topology deployments may return a 503 error. This includes, but may not be limited to, knoxsso.xml for the KnoxSSO enabled services.</p> <p>Workaround</p> <p>When this is encountered, a minor change through Ambari (whitespace even) to the knoxsso topology (or any other with this issue) and restart of the Knox gateway server should eliminate the issue.</p>
BUG-110463	<a href="#">KNOX-1434</a>	Knox	<p>Description of the problem or behavior</p> <p>Visiting Knox Admin UI in any browser (Firefox / Chrome) sets the HTTP Strict Transport Security (HSTS) header for the host where Knox is running. Any subsequent request to other service on the same host (e.g. Graphana, Ranger etc.) over HTTP would get redirected to HTTPS due to this header.</p> <p>Please note that, this HSTS header is disabled in all Knox topologies by default.</p> <p>For more information, see <a href="https://knox.apache.org/books/knox-1-1-0/user-guide.html#HTTP+Strict+Transport+Security">https://knox.apache.org/books/knox-1-1-0/user-guide.html#HTTP+Strict+Transport+Security</a></p> <p>Impact</p> <p>All the non-SSL requests to other services get redirected automatically to HTTPS and would result in SSL errors like: SSL_ERROR_RX_RECORD_TOO_LONG or some other error.</p> <p>Workaround</p> <p>Use the manager.xml topology and remove the setting from the WebAppSec provider. You can do this using the Knox Admin UI. After you have removed the setting, close your browser or clear the cookies.</p>

Hortonworks Bug ID	Apache JIRA	Apache component	Summary
BUG-101227	N/A	Spark	<p>Description of the problem or behavior</p> <p>When Spark Thriftserver has to run several queries concurrently, some of them can fail with a timeout exception when performing broadcast join.</p> <p>Associated error message</p> <pre> Caused by: java.util.concurrent.TimeoutException:   Futures timed out after [300   seconds] at  scala.concurrent.impl.Promise   \$DefaultPromise.ready(Promise.scala:219)   at scala.concurrent.impl.Promise   \$DefaultPromise.result(Promise.scala:223)   at scala.concurrent.Await\$\$anonfun   \$result\$1.apply(package.scala:107)   at scala.concurrent.BlockContext   \$DefaultBlockContext   \$.blockOn(BlockContext.scala:53) at scala.concurrent.Await   \$.result(package.scala:107) at org.apache.spark.sql.execution.joins.BroadcastH </pre> <p>Workaround</p> <p>You can resolve this issue by increasing the spark.sql.broadcastTimeout value.</p>
BUG-109979	N/A	Spark	<p>Description of the problem or behavior</p> <p>YARN NodeManagers fail to start after a Spark patch upgrade due to YarnShuffleService CNF.</p> <p>Workaround</p> <p>To resolve this problem you must:</p> <p>Replace "{{spark2_version}}" with "\${hdp.version}" in "yarn.nodemanager.aux-services.spark2_shuffle.classpath" property value. For example, old value "{{stack_root}}/{{spark2_version}}/spark2/aux/*" -&gt; new value "{{stack_root}}/\${hdp.version}/spark2/aux/*"</p>

Hortonworks Bug ID	Apache JIRA	Apache component	Summary
BUG-65977	<a href="#">SPARK-14922</a>	Spark	<p>Description of the problem or behavior</p> <p>Since Spark 2.0.0, `DROP PARTITION BY RANGE` is not supported grammatically. In other words, only `=` is supported while `&lt;`, `&gt;`, `&lt;=`, `&gt;=` aren't.</p> <p>Associated error message</p> <pre>scala&gt; sql("alter table t drop partition (b&lt;1) ").show org.apache.spark.sql.catalyst.parser.ParseException: mismatched input '&lt;' expecting '{' ','}(line 1, pos 31)</pre> <p>== SQL ==</p> <pre>alter table t drop partition (b&lt;1) -----^^^^</pre> <p>Workaround</p> <p>To drop partition, use the exact match with `=`.</p> <pre>scala&gt; sql("alter table t drop partition (b=0) ").show</pre>
BUG-106917	N/A	Sqoop	<p>Description of the problem or behavior</p> <p>In HDP 3, managed Hive tables must be transactional (<code>hive.strict.managed.tables=true</code>). Transactional tables with Parquet format are not supported by Hive. Hive imports with <code>--as-parquetfile</code> must use external tables by specifying <code>--external-table-dir</code>.</p> <p>Associated error message</p> <pre>Table db.table failed strict managed table checks due to the following reason: Table is marked as a managed table but is not transactional.</pre> <p>Workaround</p> <p>When using <code>--hive-import</code> with <code>--as-parquetfile</code>, users must also provide <code>--external-table-dir</code> with a fully qualified location of the table:</p> <pre>sqoop import ... --hive-import --as-parquetfile --external-table- dir hdfs:///path/to/table</pre>

Hortonworks Bug ID	Apache JIRA	Apache component	Summary
BUG-102672	N/A	Sqoop	<p>Description of the problem or behavior</p> <p>In HDP 3, managed Hive tables must be transactional (hive.strict.managed.tables=true). Writing transactional table with HCatalog is not supported by Hive. This leads to errors during HCatalog Sqoop imports if the specified Hive table does not exist or is not external.</p> <p>Associated error message</p> <p>Store into a transactional table db.table from Pig/Mapreduce is not supported</p> <p>Workaround</p> <p>Before running the HCatalog import with Sqoop, the user must create the external table in Hive. The --create-hcatalog-table does not support creating external tables.</p>
BUG-109607	N/A	YARN	<p>Description of the problem or behavior</p> <p>With wire encryption enabled with containerized Spark on YARN with Docker, Spark submit fails in "cluster" deployment mode. Spark submit in "client" deployment mode works successfully.</p> <p>Associated error message</p> <p>Store into a transactional table db.table from Pig/Mapreduce is not supported.</p> <p>Workaround</p> <p>There is currently no workaround.</p>



Hortonworks Bug ID	Apache JIRA	Apache component	Summary
BUG-110192	N/A	YARN	<p>Description of the problem or behavior</p> <p>When YARN is installed and configured with KNOX SSO alone, Application Timeline Server web endpoint blocks remote REST calls from YARN UI and displays a 401 Unauthorized error.</p> <p>Associated error message</p> <p>401 Unauthorized error.</p> <p>Workaround</p> <p>Administrator needs to configure Knox authentication handler for Timeline Server and existing hadoop level configuration.</p> <p>Administrator needs to tune the following cluster specific configurations. Values for the last two property is in the hadoop.authentication.* properties file.</p> <pre>&lt;property&gt; &lt;name&gt;yarn.timeline-service.http- authentication.type&lt;/name&gt; &lt;value&gt;org.apache.hadoop.security.authentication value&gt; &lt;/property&gt;  &lt;property&gt; &lt;name&gt;yarn.timeline-service.http- authentication.authentication.provider.url&lt;/ name&gt; &lt;value&gt;https://ctr- e138-1518143905142-455650-01-000002.hwx.site:44 gateway/knoxssso/api/v1/webssso&lt;/ value&gt; &lt;/property&gt;  &lt;property&gt; &lt;name&gt;yarn.timeline-service.http- authentication.public.key.pem&lt;/name&gt; &lt;value&gt;public.key.pem&lt;/value&gt; &lt;/property&gt;</pre>
RMP-11408	<a href="#">ZEPPELIN-2170</a>	Zeppelin	<p>Description of the problem or behavior</p> <p>Zeppelin does not show all WARN messages thrown by spark-shell at the Zeppelin's notebook level.</p> <p>Workaround</p> <p>There is currently no workaround for this.</p>
N/A	N/A	N/A	<p>Description of the problem or behavior</p> <p>Open JDK 8u242 is not supported as it causes Kerberos failure.</p> <p>Workaround</p> <p>Use a different version of Open JDK.</p>

**Table 3: Technical Service Bulletin**

Technical Service Bulletin	Apache JIRA	Apache component	Summary
TSB-327	<a href="#">HDFS-5698</a>	HDFS	<p>CVE-2018-11768: HDFS FSImage Corruption (potential DoS, file/dir takeover)</p> <p>In very large clusters, the in-memory format to store the user, group, acl, and extended attributes may exceed the size of the on disk format, causing corruption of fsImage.</p> <p>For the latest update on this issue, see the corresponding Knowledge article: <a href="#">CVE-2018-11768: HDFS FSImage Corruption (potential DoS, file/dir takeover)</a></p>
TSB-405	N/A	N/A	<p>Impact of LDAP Channel Binding and LDAP signing changes in Microsoft Active Directory</p> <p>Microsoft has introduced changes in LDAP Signing and LDAP Channel Binding to increase the security for communications between LDAP clients and Active Directory domain controllers. These optional changes will have an impact on how 3rd party products integrate with Active Directory using the LDAP protocol.</p> <p>Workaround</p> <p>Disable LDAP Signing and LDAP Channel Binding features in Microsoft Active Directory if they are enabled</p> <p>For the latest update on this issue see the corresponding Knowledge article: <a href="#">TSB-2021 405: Impact of LDAP Channel Binding and LDAP signing changes in Microsoft Active Directory</a></p>
TSB-406	N/A	HDFS	<p>CVE-2020-9492 Hadoop filesystem bindings (ie: webhdfs) allows credential stealing</p> <p>WebHDFS clients might send SPNEGO authorization header to remote URL without proper verification. A maliciously crafted request can trigger services to send server credentials to a webhdfs path (ie: webhdfs://...) for capturing the service principal</p> <p>For the latest update on this issue see the corresponding Knowledge article: <a href="#">TSB-2021 406: CVE-2020-9492 Hadoop filesystem bindings (ie: webhdfs) allows credential stealing</a></p>

Technical Service Bulletin	Apache JIRA	Apache component	Summary
TSB-450	<a href="#">HBASE-21000</a>	HBase	<p>Default limits for PressureAwareCompactionThroughputController are too low</p> <p>HDP and CDH releases suffer from low compaction throughput limits, which cause storefiles to back up faster than compactions can re-write them.</p> <p>For the latest update on this issue see the corresponding Knowledge article: <a href="#">Cloudera Customer Advisory: Default limits for PressureAwareCompactionThroughputController are too low</a></p>
TSB-463	N/A	HBase	<p>HBase Performance Issue</p> <p>The HDFS short-circuit setting <code>dfs.client.read.shortcircuit</code> is overwritten to disabled by <code>hbase-default.xml</code>. HDFS short-circuit reads bypass access to data in HDFS by using a domain socket (file) instead of a network socket. This alleviates the overhead of TCP to read data from HDFS which can have a meaningful improvement on HBase performance (as high as 30-40%).</p> <p>For the latest update on this issue see the corresponding Knowledge article: <a href="#">TSB 2021-463: HBase Performance Issue</a></p>
TSB-480/2	<a href="#">HIVE-24224</a>	Hive	<p>Hive ignores the property to skip a header or footer in a compressed file</p> <p>Incorrect results can occur running SELECT queries if count value is greater than 0.</p> <p>For the latest update on this issue see the corresponding Knowledge article: <a href="#">TSB 2021-480.2: Hive ignores the property to skip a header or footer in a compressed file</a></p>

Technical Service Bulletin	Apache JIRA	Apache component	Summary
TSB-494	<ul style="list-style-type: none"> <li><a href="#">HBASE-20781</a></li> <li><a href="#">HBASE-25459</a></li> <li><a href="#">PHOENIX-5250</a></li> </ul>	HBase	<p>Accumulated WAL Files Cannot be Cleaned up When Using Phoenix Secondary Global Indexes</p> <p>The Write-ahead-log (WAL) files for Phoenix tables that have secondary global indexes defined on them, cannot be automatically cleaned up by HBase, leading to excess storage usage and possible error due to filling up the storage.</p> <p>Workaround</p> <p>Perform rolling restart of HBase if the number of znodes under hbase-secure/splitWAL in ZooKeeper is greater than 8000.</p> <p>For the latest update on this issue see the corresponding Knowledge article: <a href="#">TSB 2021-494: Accumulated WAL Files Cannot be Cleaned up When Using Phoenix Secondary Global Indexes</a></p>
TSB-497	N/A	Solr	<p>CVE-2021-27905: Apache Solr SSRF vulnerability with the Replication handler</p> <p>The Apache Solr ReplicationHandler (normally registered at "/replication" under a Solr core) has a "masterUrl" (also "leaderUrl" alias) parameter. The "masterUrl" parameter is used to designate another ReplicationHandler on another Solr core to replicate index data into the local core. To help prevent the CVE-2021-27905 SSRF vulnerability, Solr should check these parameters against a similar configuration used for the "shards" parameter.</p> <p>For the latest update on this issue, see the corresponding Knowledge article: <a href="#">TSB 2021-497: CVE-2021-27905: Apache Solr SSRF vulnerability with the Replication handler</a></p>
TSB-512	N/A	HBase	<p>HBase MOB data loss</p> <p>HBase tables with the MOB feature enabled may encounter problems which result in data loss.</p> <p>For the latest update on this issue, see the corresponding Knowledge article: <a href="#">TSB 2021-512: HBase MOB data loss</a></p>

## Documentation Errata

There are no late additions or corrections to the product documentation in this release.

## Legal Information

### Hortonworks Data Platform 3.0.1

Copyright information for Hortonworks Data Platform components may be found within the documentation accompanying each component in a particular release.

Hortonworks Data Platform incorporates software from various open source projects released primarily under the Apache Software License 2.0 (“ASLv2”). Other software included may be released under the terms of alternative ASLv2 compatible open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Hortonworks Data Platform page for more information on Hortonworks Data Platform technology. For more information on Hortonworks support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Hortonworks reserves the right to change any products described herein at any time, and without notice. Hortonworks assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Hortonworks.

Trademark: Hortonworks Data Platform is/are trademark/s of Hortonworks, Inc. in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH HORTONWORKS, HORTONWORKS DOES NOT MAKE OR GIVE ANY REPRESENTATION, WARRANTY, OR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH HORTONWORKS TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. HORTONWORKS DOES NOT WARRANT THAT Hortonworks Data Platform WILL OPERATE UNINTERRUPTED OR THAT IT WILL BE FREE FROM DEFECTS OR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION OR UNAVAILABILITY, OR THAT # WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, HORTONWORKS EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.