cloudera®

Apache Accumulo Installation Guide

for using the Cloudera packaging of Accumulo for CDH

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

Release Information

Version: 1.9.2-cdh6.2.0

Date: February 22, 2019

Table of Contents

About this Guide

This guide describes how to install Cloudera's packaging of Apache Accumulo for use with CDH.

Introducing Apache Accumulo

Apache Accumulo™ is an ideal solution for government agencies looking for a secure, distributed NoSQL data store to serve their most performance-intensive Big Data applications. Accumulo is an open source project integrated with Hadoop that provides the ability to store data in massive tables (billions of rows / millions of columns) for fast, random access. Accumulo was created and contributed to the Apache Software Foundation by the National Security Agency (NSA). It has quickly gained adoption as a Hadoop-based key/value store for applications that have unique and stringent information security requirements.

Known Issues

— Unsupported features

The following upstream Apache Accumulo 1.9.2 features are not supported in the Cloudera package of Accumulo:

- User-initiated compaction strategies
- GroupBalancer
- User-specified durability
- Multi-volume deployment (including custom VolumeChoosers, including the per-table VolumeChooser).
- Data-center replication (experimental and unsupported)

For more information about problems and workarounds specific to running an Accumulo service, see the known issue document for your release of Cloudera Manager.

Prerequisites

- Accumulo depends on HDFS and ZooKeeper libraries and configuration information.
- TabletServers should be collocated with DataNodes. Optionally, you can use Accumulo with Yarn (MapReduce 2), Spark, Hive and Sqoop.
- The current release of Cloudera's package of Apache Accumulo is tested for use with CDH 6.0.0 and higher.
- Running Apache Accumulo 1.9.2 is supported on the following OSes with the same minor versions that are supported for the CDH installation with which Apache Accumulo is installed:
  - RHEL 6 / CentOS 6 / OL 6
  - RHEL 7 / CentOS 7 / OL 7

- ○ SLES 12
  - ○ Ubuntu 16 Xenial Xerus
- Running Apache Accumulo 1.9.2 is supported on Java 8 versions that are supported for the CDH install with which Apache Accumulo is installed.
- For full cluster installations, Cloudera strongly recommends following the Tips and Guidelines in the CDH 6 Installation Guide.

Install Apache Accumulo Using Cloudera Manager

The process of installing Cloudera's package of Accumulo by using Cloudera Manager 6.0.0 or higher is the following:

1. Install Accumulo CSD
2. Install CDH and Accumulo Parcel
3. Configure HDFS
4. Add the Accumulo Service
5. Optionally: Performance Tuning - Relaxing WAL Durability
6. Optional: Security Configuration - Kerberos for Accumulo Clients

> **Note:** HDFS and Zookeeper are mandatory services for Accumulo. Other services are optional.

Step 1: Install Accumulo CSD

Integrating an add-on service requires a Custom Service Descriptor (CSD) file. A CSD file contains all the configuration needed to describe and manage a new service. A CSD is provided in the form of a JAR file. For more information about CSD file, see Custom Service Descriptor File.

> **Important:** If you use CM 6.2.0 or higher, skip Step 1 and continue the procedure from Step 2.

1. Download the jar file for the CSD that provides the Accumulo-on-C6 service: ACCUMULO_C6-6.1.0.jar.[1]
2. Copy **ACCUMULO_C6-6.1.0.jar** to the location Cloudera Manager is configured to look for CSD files. That is by default **/opt/cloudera/csd**, on the host where Cloudera Manager is running.
3. Restart Cloudera Manager as described in Configuring the Location of Customer Service.

---

[1]https://archive.cloudera.com/p/accumulo6/6.1.0/csd/ACCUMULO_C6-6.1.0.jar

Step 2: Install CDH and Accumulo Parcels

1. Ensure that the **Use Parcels** option is checked.

Install Method
- ○ Use Packages
- ● Use Parcels (Recommended) | More Options | Proxy Settings

2. Select version **CDH-6.0.0-1.cdh6.0.0.p0.537114** or higher for the CDH version.

CDH Version
- ● CDH-6.0.0-1.cdh6.0.0.p0.537114
- ○ CDH-5.14.5-1.cdh5.14.5.p0.101

3. Select version **ACCUMULO-1.9.2-1.ACCUMULO6.1.0.p0.908695** or higher for the Accumulo parcel.

Additional Parcels
- ● ACCUMULO-1.9.2-1.ACCUMULO6.1.0.p0.908695
- ○ ACCUMULO-1.7.2-5.5.0.ACCUMULO5.5.0.p0.8
- ○ None

   If an Accumulo parcel is not available, do the following:
   a. Set up the parcel URL under **Host/Parcel/Configuration** by adding the URL https://archive.cloudera.com/p/accumulo6/6.1.0/parcels/ to the **Remote Parcel Repository URLs** list.
   b. Save your changes.
4. Click **Continue**.

Step 3: Configure HDFS

Cloudera strongly recommends that you establish an HDFS Name Service on the cluster that will run Accumulo. Due to the way Accumulo manages files within HDFS, doing so greatly reduces administrative tasks in the future if a NameNode is replaced or moved. To set up an HDFS Name Service, see Enabling HDFS High Availability.

To prevent data loss, configure HDFS to durably write data on file close. If the following configuration changes are not made, Accumulo issues warning messages until the problem is corrected.

1. Navigate to your cluster's **HDFS service** page.

2. Click the **Configuration** tab.

3. Search for "*hdfs-site.xml*".

4. Find the HDFS-1 (Service-Wide) properties for **HDFS Service Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml**.

5. Set the following:

Name: dfs.datanode.synconclose

Value: true

6. Find any Gateway group properties for **HDFS Client Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml**.

7. Set the following:

     Name: dfs.datanode.synconclose

     Value: true

8. Save your changes.

9. **Restart Stale Services** with Re-deploy client configuration.

> **Note:** If you want to run Accumulo with HDFS encryption, you must disable HDFS trash, which Accumulo uses by default. HDFS trash is not compatible with HDFS Transparent Encryption. To disable HDFS trash in Accumulo, see HDFS storage demands due to retained HDFS trash. For information about HDFS encryption, see HDFS Transparent Encryption.

Step 4: Add the Accumulo Service

1. Navigate to the **Cloudera Manager Admin Console** home page.

2. Click the actions menu for the cluster you want to add the Accumulo service to (for example, Cluster 1), and select **Add Service**.
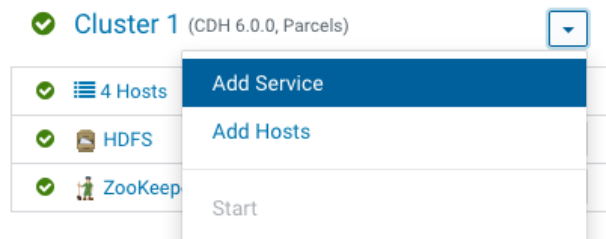
3. Select **Accumulo** and click **Continue**.

4. Select the dependent services and click **Continue**.

5. Assign the Accumulo roles to the hosts in your cluster.

     - Assign a **Tablet Server** role on each host that is assigned the **DataNode** role.

     - Assign the **Monitor**, **Garbage Collector**, **Tracer**, and **Master** roles to **non-DataNodes**.

     - Assign the **Gateway** role to any hosts where you want to use Accumulo and that do not already have other Accumulo roles assigned.

6. Click **Continue**.

7. Configure the **Accumulo Instance Secret**.

> **Important:** It is strongly recommended to protect this secret, because the security of the
> Accumulo service relies on it.

8. Set the **Accumulo Instance Name**.

9. Set the **Trace User** and **Trace Password** settings.

   It is strongly recommended that you not leave the **Trace User** set to **root**, because this is not a
   secure configuration.

10. If this is the first installation of Accumulo, ensure that the **Initialize Accumulo** is checked.

    If you are upgrading Accumulo, see Upgrading from CDH 5 to CDH 6.

11. Click **Continue**.

12. Wait while Cloudera Manager does the initial service set up.

13. Click **Continue**.

14. Click **Finish**.

15. Verify your installation using Test the Accumulo Shell.

Step 5: Optional: Performance Tuning – Relaxing WAL Durability

For increased write throughput, use the `BatchWriter API` to ingest data into Accumulo. However,
using this API reduces data durability. Use this setting only in environments with reliable UPS.

To enable this setting, perform the following configuration changes:

1. Navigate to your cluster's **Accumulo** service page.

2. Click **Configuration** and then click **View and Edit**.

3. Search for "*Tablet Server accumulo-site.xml*".

4. Find the the Tablet Server Default Group / Advanced section property for **Tablet Server
   Advanced Configuration Snippet (Safety Valve) for accumulo-site.xml**.

5. Set the following:

   Name: table.durability

   Value: flush

6. Save your changes.

7. **Restart Stale Services**.

For more information about durability performance, see [Durability Performance Implications](#).

Step 6: Optional: Security Configuration - Kerberos for Accumulo Clients

If you enable Kerberos for your CDH cluster, Cloudera Manager makes the configuration changes required for the Accumulo service to run top of secured HDFS and ZooKeeper services.

Additionally, you can enable Kerberos-based authentication and authorization for clients of the Accumulo cluster by making the following changes.

> **Note**: These changes require cluster downtime period.

1. Go to the **Accumulo** service page of your cluster.
2. Click **Configuration** and then click **View and Edit**.
3. Find the Accumulo (Service-wide) / Main section properties for **Trace User**.
4. Click the field and set the value to the principal you will use for the trace user.

   It can be with an instance, for example `trace/some.host.example.com@EXAMPLE.COM` or without an instance, for example, `trace@EXAMPLE.COM`.

5. Find the Accumulo (Service-wide)/Advanced section property for **Service Advanced Configuration Snippet (Safety Valve) for accumulo-site.xml**.
6. Click the field and add the following snippet:

```
<property>
      <name>instance.rpc.sasl.enabled</name>
      <value>true</value>
</property>

<property>
      <name>rpc.sasl.qop</name>
      <value>authconf</value>
      <description>
      one of "auth", "auth-int", or "auth-conf". If you use delegation
      tokens you must set this to "auth-conf" to avoid leaking tokens
      to network observers.
      </description>
</property>

<property>
      <name>instance.security.authenticator</name>
      <value>org.apache.accumulo.server.security.handler.KerberosAuthe
```

```
        nticator</value>
</property>

<property>
        <name>instance.security.authorizor</name>
        <value>org.apache.accumulo.server.security.handler.KerberosAutho
        rizor</value>
</property>

<property>
        <name>instance.security.permissionHandler</name>
        <value>org.apache.accumulo.server.security.handler.KerberosPermi
        ssionHandler </value>
</property>

<property>
        <name>trace.token.type</name>
        <value>org.apache.accumulo.core.client.security.tokens.KerberosT
        oken</value>
</property>

<property>
        <name>trace.token.property.keytab</name>
        <value>accumulo_c6.keytab</value>
        <description>
        Path to a keytab corresponding to the principal given in
        trace.user. Will be used by the Trace and Monitor roles to write
        and read respectively the trace table.
        </description>
</property>
```

7. Save your changes.

8. Stop all Tracer roles in the cluster.

9. **Restart Stale Services** with [Re-deploy client configuration](#).

10. Restart all roles except Tracers.

11. On each Gateway machine where you will run an Accumulo client, create a client configuration
    file for each user who will run an Accumulo client. Create that file in `~/.accumulo/config`
    with the following contents:

```
# Each of these must match what's in the server configs exactly
instance.rpc.sasl.enabled=true
```

```
rpc.sasl.qop=auth/conf

# Should match the System User configuration value
kerberos.server.primary=accumulo
```

12. Set up an administrative user using the [Administrative User section of he user manual](#).

> **Note**: The remaining steps in the Security Configuration - Kerberos for Accumulo Clients process presume this user is named `accumulo_admin@EXAMPLE.COM`.

13. Set up access to the trace table for the trace principal.

    For example, on a Gateway machine with the client configuration listed in Step 11, log in to Kerberos using the administrative user from Step 12, and then run the following Accumulo shell commands:

```
$ accumulo shell
Shell - Apache Accumulo Interactive Shell
-
- version: 1.9.2-cdh6.1.0
- instance name: cloudera
- instance id: e30831be-b22f-4dfb-99e3-40640c5ae82f
-
- type 'help' for a list of available commands
-
root@accumulo> createuser trace@EXAMPLE.COM
root@accumulo> grant Table.READ -t trace -u trace@EXAMPLE.COM
root@accumulo> grant Table.WRITE -t trace -u trace@EXAMPLE.COM
root@accumulo> grant Table.ALTER -t trace -u trace@EXAMPLE.COM
```

    If you configure the Trace User with a `_HOST` instance component, you must repeat the previous command with an instance for each host that will run a Tracer or Monitor role.

14. Start all Tracer roles.

Upgrading from CDH 5 to CDH 6

Before upgrading, plan for downtime. Cloudera does not support rolling upgrades from CDH 5 to CDH 6. Cloudera supports only upgrading from Accumulo 1.7 version to Accumulo 1.9. Additionally, you must remove and reinstall the Accumulo service as a part of upgrading.

Step 1: Document and Remove Accumulo 1.7 for CDH 5

1. Ensure there are no ongoing administrative operations for the Accumulo service.

2. Check if there are any FATE transactions in Accumulo shell using:

```
root@accumulo> fate print
```

3. Remove them using:

    ```
    root@accumulo> fate delete <transaction id>
    ```

4. Stop the Accumulo 1.7 service.

5. Go to the Accumulo 1.7 service page.

6. Click the **Configuration** tab.

7. Click through the configuration settings and make a note of any non-default settings. You will
    need this information in Step 3: Add Accumulo 1.9 for CDH 6.

    You can use the Cloudera Manager REST API to export and import all of its configuration data.
    For more information, see Backing Up and Restoring the Cloudera Manager Configuration.

8. Click the **Instances** tab.

9. Make a note of your role assignments. You will need this information in Step 3: Add Accumulo
    1.9 for CDH 6.

10. Return to the **Cloudera Manager Admin Console** home page.

11. Under actions for the Accumulo 1.7 Service select **Delete**.

Step 2: Update Parcels and Upgrade CDH

1. Follow the instruction described in Upgrading the CDH Cluster and upgrade your cluster from
    CDH 5 to CDH 6.

> **Note:** Cloudera Manager will ask you to remove Accumulo if you have not done it yet.
>
> ⌄ ❌ Configuration Check
>
> > ❌ The cluster is currently running an Accumulo service, which is not supported in CDH 6.0.0-1.cdh6.0.0.p0.537114. You cannot upgrade to CDH
> >   6.0.0-1.cdh6.0.0.p0.537114 until you remove the ACCUMULO16-1 service from the cluster.
>
> All Configuration Issues

2. Install the Accumulo CSD as described in Install Accumulo CSD.

3. Go to the **Hosts** and then the **Parcels** page.

Parcel Name	Version	Status	
ACCUMULO	1.9.2-1.ACCUMULO6.0.0.p0.645027	Available Remotely	Download
	1.7.2-5.5.0.ACCUMULO5.5.0.p0.7	Distributed	Activate ▾

4. You must **Deactivate** the active parcel for **1.7.2-5.5.0.ACCUMULO5.5.0** or similar.

5. You must **Remove from Hosts** and then **Delete** this parcel.

6. **Download the** Accumulo **1.9.2-1.ACCUMULO6.1.0** parcel.

7. **Distribute** and **Activate** the parcel.

Step 3: Add Accumulo 1.9 for CDH 6

1. Add the Accumulo 1.9 service to your cluster.
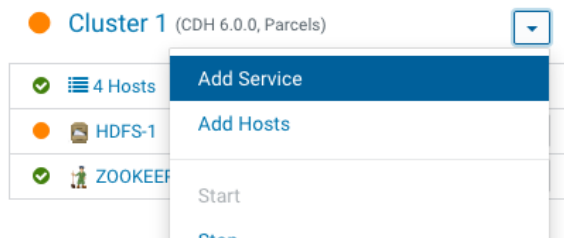
2. Assign roles to match those you wrote down in Step 1: Document and Remove Accumulo 1.7 for CDH 5.

3. In the initial configuration, ensure the **Instance Name** and **Secret** match those you wrote down in Step 1: Document and Remove Accumulo 1.7 for CDH 5.

4. Uncheck **Initialize Accumulo** checkbox.

5. Click **Continue** and finish the installation.

   The Accumulo Service is now up and running.

Troubleshooting the Accumulo Installation

Under‑replicated Block Exceptions or Cluster Failure Occurs on Small Clusters

By default Accumulo attempts to use a replication factor of five for the HDFS files that make up the metadata table. It ignores the `table.file.replication` setting. This causes under-replication warnings until you correct the number of nodes or until you manually adjust the replication setting on that table.

Normally, Cloudera Manager does not set a maximum replication factor. If it is set, Accumulo use the `table.file.replication` setting only on normal user tables to that maximum.

However, if `dfs.replication.max` is adjusted to match the number of cluster nodes, attempts by Accumulo to create new files for its internal tables will fail on clusters with fewer than five datanodes.

Do the following to fix the issue:

1. Edit `dfs.replication.max` setting for HDFS to be >= 5.
2. Adjust replication on the metadata and root tables to be less than or equal to the number of DataNodes.

3. Readjust `dfs.replication.max` to lower it again.

For example, to adjust the replication in the Accumulo shell:

```
root@accumulo> config -t accumulo.metadata -s table.file.replication=3
root@accumulo> config -t accumulo.root -s table.file.replication=3
```

HDFS Storage Demands Due to Retained HDFS Trash

By default, Accumulo uses [HDFS trash](#) if it is enabled for all files it deletes, including write-ahead logs and long-term storage files that were blocked due to compaction. By default, the retention period for the HDFS trash is 24 hours. On Accumulo installations with a heavy write workload, this can result in a large amount of data accumulating in the trash folder for the service user.

As a workaround, periodically run the `hdfs dfs -expunge` command as the Accumulo service user. The command must be run twice each time you want to purge a backlog of data. The first time creates a trash checkpoint, and the second removes that checkpoint immediately.

Alternatively, you can tune the amount of time HDFS retains trash to control how much data Accumulo saves. This change is HDFS-wide and impacts the ability to recover from accidental deletions unrelated to Accumulo.

Do the following to change the HDFS trash setting in Cloudera Manager:

1. Go to the **HDFS service** page.
2. Click the **Configuration** tab.
3. Search for "*trash*".
4. Change the **Filesystem Trash Interval** to a smaller value. For example, 4 hours.
5. Save your changes.
6. **Restart Stale Services** with [Re-deploy client configuration](#).

In some deployments, you might not want to change the system-wide retention period for HDFS trash. If that is the case, you can disable Accumulo's use of the HDFS trash entirely. If you do so, any deletions through the Accumulo APIs are unrecoverable.

Do the following to configure Accumulo to skip the HDFS trash in Cloudera Manager:

1. Go to the **Accumulo service** page.
2. Click the **Configuration** tab.
3. Search for "*accumulo-site.xml*".
4. Find the Accumulo (Service-wide) / Advanced section's property for **Accumulo Service Advanced Configuration Snippet (Safety Valve) for accumulo-site.xml**.
5. Set the following:

    Name: gc.trash.ignore

Value: true

6. Save your changes.

7. **Restart Stale Services**.

Test the Accumulo Shell

You can run the Accumulo shell on server hosts and hosts that have the Gateway role in your cluster.

By default, the user **root** is created and given the password **secret**. If you did not set a different password during installation, Cloudera strongly recommends that you change the root user password.

The following steps verify that the Accumulo shell is working and allows you to change the root user password.

1. Launch the Accumulo shell for the default root user.

```
$ accumulo shell -u root

Password: ******

Shell - Apache Accumulo Interactive Shell

-

- version: 1.9.2-cdh6.1.0

- instance name: accumulo

- instance id: 9863d1f1-c323-4671-9e1f-69857f0d635f

-

- type 'help' for a list of available commands

-

root@accumulo>
```

2. Use the `passwd` command to set a new password for the root user.

```
root@accumulo> passwd

Enter current password for 'root': ******

Enter new password for 'root': *********

Please confirm new password for 'root': *********

root@accumulo>
```

3. Relaunch the shell with this new password.

```
root@accumulo> exit
```

```
$ accumulo shell -u root

Password: ******

Shell - Apache Accumulo Interactive Shell

-

- version: 1.9.2-cdh6.1.0

- instance name: accumulo

- instance id: 9863d1f1-c323-4671-9e1f-69857f0d635f

-

- type 'help' for a list of available commands

-

root@accumulo>
```

4. Verify that you can list tables.

```
root@accumulo> tables

accumulo.metadata

accumulo.replication

accumulo.root

trace

root@accumulo>
```

5. If the trace table does not exist, make sure that you have created the trace user. Use the same password you used for the **trace.token.property.password** setting in the **Trace Password** setting in Cloudera Manager.

```
root@accumulo> createuser trace

Enter new password for 'trace': *****

Please confirm new password for 'trace': *****

root@accumulo> grant System.CREATE_TABLE -s -u trace

root@accumulo> tables

accumulo.metadata

accumulo.replication

accumulo.root

trace
```

```
root@accumulo> revoke System.CREATE_TABLE -s -u trace
```

For more information on using the Accumulo shell, see the [Accumulo user manual.](#)

Using Sqoop with Accumulo

CDH 6.0.0 and higher include Sqoop bindings for import/export of data with Accumulo. For instructions on invoking Sqoop with Accumulo as a source or sink, see the [Sqoop documentation.](#)

When running the Sqoop command, you may see warning messages about failing to create `/usr/lib/accumulo/logs`. These messages are safe to ignore.

To use Sqoop integration, you must perform the following configuration changes:

1. Navigate to your cluster's Sqoop Client service page.

2. Click **Configuration**.

3. Search for "*sqoop-env.sh*".

4. Find the Gateway Default Group / Advanced section's property for **Sqoop Client Client Advanced Configuration Snippet (Safety Valve) for sqoop-conf/sqoop-env.sh.**

5. Click the field and add the snippet appropriate for your installation, ensuring that each line ends with a bash comment hash, '#'.

```
export ACCUMULO_CONF_DIR=/etc/accumulo/conf #

export ACCUMULO_HOME=/opt/cloudera/parcels/ACCUMULO/lib/accumulo #

export HADOOP_CLIENT_HOME=/opt/cloudera/parcels/CDH/lib/hadoop/client #

export HADOOP_PREFIX=/opt/cloudera/parcels/CDH/lib/hadoop #

export ZOOKEEPER_HOME=/opt/cloudera/parcels/CDH/lib/zookeeper #
```

6. Save your changes

7. **Restart Stale Services** with [Re-deploy client configuration](#).

Using LZO Compression with Accumulo

Optionally, you may enable LZO compression for use with Accumulo.

If you install it based on distribution packaging, no special instructions are needed.

If you are using parcels for the installation, do the following:

1. Install the GPLEXTRAS Parcel.

2. Go to the HDFS service.

3. Select the **Configuration** tab.

4. Search for the **io.compression.codecs** property.

5. In the **Compression Codecs** entry, click in the field and click the + to add a new entry.

6. Add the following codecs for LZO:

   com.hadoop.compression.lzo.LzoCodec

   com.hadoop.compression.lzo.LzopCodec

7. Save the configuration for the HDFS Service.

8. Go to the Accumulo service.

9. Select the **Configuration** tab.

10. Search for the *"accumulo_classpath"* property.

11. In the **Additional Classpath** entry, click in the field (if needed, click + to add a new entry).

12. Add the following path that corresponds to the version you are using

    `/opt/cloudera/parcels/GPLEXTRAS/lib/hadoop/lib/hadoop-lzo.jar`

    > **Note:** If you have configured Cloudera Manager to install parcels in a location other than `/opt/cloudera/parcels/`, change the above paths to reflect that location.

13. Search for *"service environment"*.

14. Find the Service Wide / Advanced section's **Accumulo Service Environment Advanced Configuration Snippet (Safety Valve)**.

15. Click the field and add the following line that corresponds to the version you are using:

    `LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/opt/cloudera/parcels/GPLEXTRAS/lib/hadoop/lib/native`

    > **Note:** If you have configured Cloudera Manager to install parcels in a location other than `/opt/cloudera/parcels/`, change the above paths to reflect that location.

16. Search for the *"accumulo_java_opts"* property.

17. In the **Additional Java Options** entry, click the field.

18. Add to the end of the existing options:

    `-Dio.compression.codec.lzo.class=com.hadoop.compression.lzo.LzoCodec`

19. Save the Accumulo service.

20. **Restart Stale Services**.

Using Accumulo with Maven

If you want to build applications or tools with Cloudera's packaging of Accumulo and you are using Maven or Ivy for dependency management, you can pull the Accumulo artifacts from the Cloudera Maven repository. The repository is available at https://repository.cloudera.com/artifactory/cloudera-repos/. The following is a sample snippet from a POM (*pom.xml*) file:

```
<repositories>
  <repository>
    <id>cloudera</id>
    <name>Cloudera Releases Repository</name>
    <url>https://repository.cloudera.com/artifactory/cloudera-repos/</url>
  </repository>
</repositories>
```

CDH 6-Compatible Releases

The following table lists the project name, groupId, artifactId, and version required to access each CDH 6-compatible artifact. Client applications normally only requires the `accumulo-core` artifact as a dependency and may need the `accumulo-maven-plugin` for running integration tests.

Project	groupId	artifactId	version
Accumulo	org.apache.accumulo	accumulo-core	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-examples-simple	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-fate	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-gc	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-master	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-maven-plugin	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-minicluster	1.9.2-cdh6.1.0

	org.apache.accumulo	accumulo-monitor	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-native	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-project	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-proxy	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-server-base	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-shell	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-start	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-test	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-trace	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-tracer	1.9.2-cdh6.1.0
	org.apache.accumulo	accumulo-tserver	1.9.2-cdh6.1.0

Default Ports

If your cluster is running firewall software, you may need to allow communication between hosts on specific ports. The following table lists the default port for each server process and the configuration property used to change that value.

Accumulo Process	Port	Property
Garbage Collector	50091	gc.port.client
Master	10010	master.port.client
Monitor (Log Forwarding)	4560	monitor.port.log4j
Monitor (Client Port)	50095	monitor.port.client

Tablet Server	10011	tserver.port.client
Tracer	12234	trace.port.client
Master Replication Service	10001	master.replication.coordinator.port
TabletServer Replication Service	10002	replication.receipt.service.port