

Operational Database powered by Apache Accumulo 2.1.2

Installing OpDB powered by Apache Accumulo

Date published: 2021-03-03

Date modified: 2024-10-30

CLOUDERA

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

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

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks 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 CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA 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.

Contents

Installing Accumulo Parcel 2.1.2.....	4
Install Operational Database powered by Apache Accumulo.....	4
Install Operational Database powered by Apache Accumulo CSD file.....	4
Install CDP.....	4
Install Operational Database powered by Apache Accumulo parcel.....	5
Add Accumulo on CDP service.....	6
Verify your Operational Database powered by Apache Accumulo installation.....	9
Installing Accumulo Parcel 1.1.0.....	11
Install Operational Database powered by Apache Accumulo.....	11
Install Operational Database powered by Apache Accumulo CSD file.....	11
Install CDP.....	12
Install Operational Database powered by Apache Accumulo parcel.....	12
Add Accumulo on CDP service.....	15
Verify your Operational Database powered by Apache Accumulo installation.....	18
Installing Accumulo Parcel 1.10.....	20
Install Accumulo.....	20
Install Accumulo CSD file.....	20
Install CDP.....	20
Install Accumulo 1.10 parcel.....	21
Add Accumulo on CDP service.....	24
Customizing Kerberos Principals.....	27
Creating a trace user in unsecure Accumulo deployment.....	29
Check trace table.....	30
Provide user permissions.....	31
Verify your Accumulo installation.....	32
Upgrading Accumulo from 1.10.3 to 2.1.2.....	34
Removing Accumulo 1.10.3.....	34
Removing and updating Accumulo parcels.....	35
Adding Accumulo service (unsecure).....	35
Adding Accumulo service (secure).....	36
Upgrading Accumulo from 1.1.0 to 2.1.2.....	37
Removing Accumulo 1.1.0.....	37
Removing and updating Accumulo parcels.....	38
Adding Accumulo service (unsecure).....	39
Adding Accumulo service (secure).....	39

Installing Accumulo Parcel 2.1.2

Learn how to install the Operational Database powered by Apache Accumulo (Accumulo-on-cdp 2.1.2 parcel) on your CDP Private Cloud Base environment.

Install Operational Database powered by Apache Accumulo

Before installing Operational Database powered by Apache Accumulo (OpDB), you first have to install Cloudera Manager, CDP cluster services, then install the OpDB CSD and the OpDB Parcel. Ensure you verify your installation by running certain basic commands.

Install Operational Database powered by Apache Accumulo CSD file

As Operational Database powered by Apache Accumulo (OpDB) is an add-on service, you have to install its Custom Services Descriptor (CSD) file to manage the Accumulo on CDP service.

About this task

A Custom Services Descriptor (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 files, see the *Custom Service Descriptor Files* section in the *Add-on Services* documentation.

Before you begin

Install Cloudera Manager.

Procedure

1. Download the ACCUMULO_CDP-2.1.2_7.1.9.0.0.jar jar file for the CSD that provides the Accumulo on CDP service.

You can download that JAR file from:

```
archive.cloudera.com/p/accumulo7/2.1.2_7.1.9.0.0/csd/
```

2. Copy the ACCUMULO_CDP-2.1.2_7.1.9.0.0.jar file to the location where Cloudera Manager is configured to look for the CSD files. That is by default in the /opt/cloudera/csd folder on the host where Cloudera Manager is running.
3. Restart Cloudera Manager as described in the *Configuring the Location of Customer Service* section in the *Add-on Services* documentation.

What to do next

Install CDP cluster and services.

Install CDP

Before you install the Operational Database powered by Apache Accumulo (OpDB) parcel, you have to install CDP Private Cloud Base and add the mandatory services.

The OpDB service is only supported on CDP Private Cloud Base. For full cluster installation Cloudera recommends following the [CDP Private Cloud Base Installation Guide](#).

For a FIPS compliant OpDB service you have to follow the FIPS deployment CDP setup steps. For more information, see [Installing and Configuring CDP with FIPS](#).

Ensure you install the following mandatory services for OpDB powered by Apache Accumulo:

- HDFS

- ZooKeeper

Install Operational Database powered by Apache Accumulo parcel

After you have installed the Operational Database powered by Apache Accumulo (OpDB) CSD and the required CDP cluster and services, you have to install the OpDB parcel.

There are two options to install the OpDB parcel:

- Using a Local Parcel Repository
- Using Remote Parcel Repository URLs (archive.cloudera.com)

Install Operational Database powered by Apache Accumulo parcel using Local Parcel Repository

You can install the OpDB powered by Apache Accumulo Parcel (OpDB) using your Local Parcel Repository. For this, you have to download the parcel, copy it to your Local Parcel Repository, and then distribute and activate it using Cloudera Manager.

Before you begin

- Install the OpDB CSD file
- Install CDP and add the HDFS and ZooKeeper services to your deployment.

Procedure

1. Download the ACCUMULO_ON_CDP-2.1.2_7.1.9.0.0.p0.11738502-el7.parcel OpDB parcel from:

```
archive.cloudera.com/p/accumulo7/2.1.2_7.1.9.0.0/parcels/
```

2. Copy the ACCUMULO_ON_CDP-2.1.2_7.1.9.0.0.p0.11738502-el7.parcel file to the Local Parcel Repository. The Local Parcel Repository is by default located in /opt/cloudera/parcel-repo directory on the host where Cloudera Manager is running.
3. In Cloudera Manager, navigate to Parcels.
4. Click Check for New Parcels. The ACCUMULO_ON_CDP-2.1.2_7.1.9.0.0.p0.11738502-el7.parcel parcel appears in the Parcel list.
5. In the row of ACCUMULO_ON_CDP-2.1.2_7.1.9.0.0.p0.11738502-el7.parcel, click Distribute.
6. After the distribution is completed, click Activate.

Results

OpDB parcel is installed.

What to do next

Add the Accumulo on CDP service.

Install Operational Database powered by Apache Accumulo using Remote Parcel Repository


You can install the Operational Database powered by Apache Accumulo (OpDB) parcel using the archive.cloudera.com Remote Parcel Repository.

Before you begin

- Install the OpDB CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.

Procedure

1. In Cloudera Manager, navigate to Parcels.
2. Click Parcel Repositories & Network Settings.

3. In the Remote Parcel Repository URLs section, click  (Add Another) to include another remote parcel repository URL.
4. Add the following URL:

```
archive.cloudera.com/p/accumulo7/2.1.2_7.1.9.0.0/parcels/
```
5. Click Save & Verify Configuration.
6. In the row of ACCUMULO_ON_CDP-2.1.2_7.1.9.0.0.p0.11738502-el7.parcel, click Download.
7. After the download is completed, click Distribute.
8. Once the distribution is completed, click Activate.

What to do next

Add the Accumulo on CDP service.

Add Accumulo on CDP service

After you have installed the Operational Database powered by Apache Accumulo (OpDB) CSD, the required CDP cluster and services, and the OpDB parcel, you can add the Accumulo on CDP service to your cluster.

You can add a secure or an unsecured Accumulo on CDP service to your CDP Private Cloud Base deployment when installing OpDB. A secure Accumulo on CDP service requires a secure cluster, which means an SASL-enabled Kerberized cluster.

The steps to add a secure Accumulo on CDP service are different from the steps to add an unsecure Accumulo on CDP service. If you require a FIPS compliant OpDB service, follow the steps in the Add secure Accumulo on CDP service to your cluster.

Add unsecure Accumulo on CDP service to your cluster

If your cluster is not an SASL-enabled Kerberized cluster, add an unsecure Accumulo on CDP service to your cluster when installing an Operational Database powered by Apache Accumulo (OpDB).

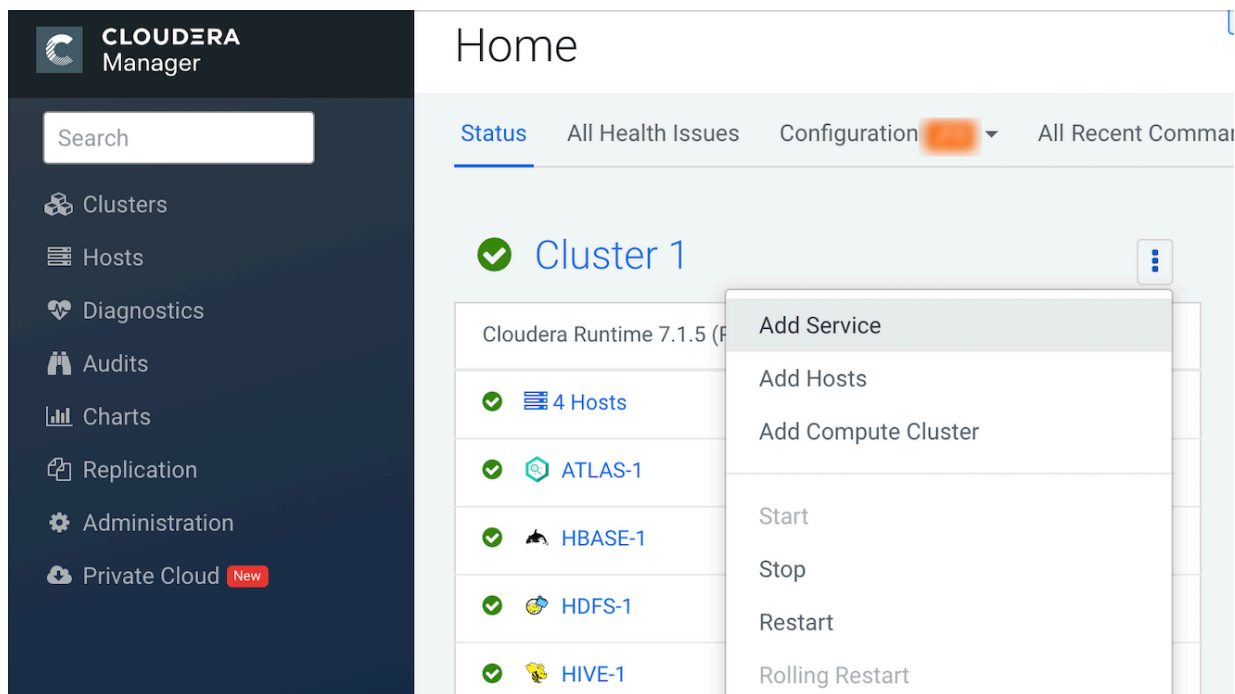
Before you begin

- Install the OpDB CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.
- Install the OpDB parcel.

Procedure

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

- Click the action menu of the cluster to which you want to add the Accumulo on CDP service (for example, Cluster 1), and select Add Service.



- Select Accumulo on CDP and click Continue.
- Select the dependent services and click Continue.
- Assign the Accumulo roles to the hosts in your cluster:
 - Assign a Table Server role on each host to which the DataNode role is assigned.
 - Assign the Monitor, Garbage Collector, and Master role to non-DataNodes.
 - Assign the Gateway role to any hosts where you want to use the Accumulo on CDP service, and where you do not already have any other OpDB roles assigned.
- Click Continue.
- Configure the Accumulo Instance Secret.



Important: Cloudera recommends to protect this information, because the security of the Accumulo service relies on it.

- Set the Accumulo Instance Name.
- If this is the first installation of OpDB, ensure that the Initial Accumulo option is selected.

Instance.volume specifies the filesystem URIs where Accumulo should store its data. It should be set when the you want to explicitly control where Accumulo stores its data. If you do not explicitly set the instance.volumes option, Accumulo will automatically use the default filesystem URI from the Hadoop configuration.
- Click Continue.

Wait while Cloudera Manager does the initial service setup.
- After the initial service setup is complete, click Continue.
- Click Finish.

Add secure Accumulo on CDP service to your cluster

To ensure data availability to diverse sets of users with varying levels of permissions and security clearance, add a secure Accumulo on CDP service.

About this task

If your secure Accumulo on CDP service needs to be FIPS compliant, ensure that you have followed the FIPS specific deployment setup steps provided in [Installing and Configuring CDP with FIPS](#).

Before you begin

- Install OpDB CSD file.
- Install CDP.
- Add the HDFS and ZooKeeper services to your CDP deployment.
- Install the OpDB parcel.
- Ensure that you have a SASL-enabled Kerberized cluster.

Procedure

1. Navigate to the Cloudera Manager Admin Console home page.
2. Select the action menu of the cluster to which you want to add the Accumulo on CDP service (for example, Cluster 1), and select Add Service.

3. Select Accumulo on CDP and click Continue.
4. Select the dependent services and click Continue.
5. Assign the Accumulo roles to the hosts in your cluster:
 - Assign a Table Server role on each host to which the DataNode role is assigned.
 - Assign the Monitor, Garbage Collector, and Master role to non-DataNodes.
 - Assign the Gateway role to any hosts here you want to use the OpDB service and that do not already have other OpDB roles assigned.
6. Click Continue.
7. Configure the Accumulo Instance Secret.



Important: Cloudera recommends to protect this information, because the security of the Accumulo service relies on it.

8. Set the Accumulo Instance Name.

9. Select Enable Kerberos Authentication.

Create a client configuration file for the admin user.

This step needs to be performed on each Gateway machine where you will run an OpDB client.

Client configuration can be created by editing the configuration file in the following way:

```
auth.principal=accumulo/[***HOSTNAME***]
auth.type=kerberos
auth.token=/var/run/cloudera-scm-agent/process/[***TIMESTAMP***]-[***SERVICE NAME***]-[***INSTANCE NAME***]/accumulo_on_cdp.keytab
sasl.enabled=true
sasl.qop=auth-conf
```

By default the Accumulo principal on which master is running becomes the administrator principal. You can set up a different administrative user. For more information, see the [Administrative User section of the user manual](#). In this case, you need to change the `auth.principal` and `auth.token` configuration properties accordingly.

10. If this is the first installation of OpDB, ensure that the Initial Accumulo option is selected.**11. Specify Instance.volume.**

This parameter specifies the filesystem URIs where Accumulo stores its data. It must be set when you want to explicitly control where Accumulo stores its data. If you do not explicitly set the `instance.volumes` option, Accumulo automatically uses the default filesystem URI from the Hadoop configuration.

12. Click Continue.

Wait while Cloudera Manager does the initial service setup.

13. After the initial service setup is complete, click Continue.**14. Click Finish.****Related Information**

[Customizing Kerberos Principals](#)

Verify your Operational Database powered by Apache Accumulo installation

You can verify your Operational Database powered by Apache Accumulo (OpDB) installation using the Accumulo shell.

Before you begin

- Install the OpDB CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.
- Install the OpDB parcel.
- Add the Accumulo on CDP service to your deployment.

Procedure**1. Launch the Accumulo shell for the admin user.**

- Unsecure OpDB

```
$ accumulo shell -u root
Password: *****
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
```

```
root@accumulo>
```

- Secure OpDB

```
kinit -kt /var/run/cloudera-scm-agent/process/[***TIMESTAMP***]-[***SERVICE NAME***]-[***INSTANCE NAME***]/accumulo_on_cdp.keytab accumulo/${***HOSTNAME***}
accumulo shell
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
accumulo/hostname@REALM.COM@accumulo>
```

2. Run the following commands to verify your installation:

- Create a table: createtable test
The test table is created and becomes the current table.
- List the admin user's permission and ensure it has admin and system level privileges: userpermissions
- On a secure cluster, the user who creates the table does not get access to the new table by default. These permissions have to be set manually:

```
'grant Table.READ -t test -u accumulo/hostname@REALM.COM'
'grant Table.WRITE -t test -u accumulo/hostname@REALM.COM'
'grant Table.ALTER_TABLE -t test -u accumulo/hostname@REALM.COM'
'grant Table.DROP_TABLE -t test -u accumulo/hostname@REALM.COM'
```

- Insert some rows to the test table:

```
insert row f1 q1 v1
insert row f2 q2 v2
insert row f3 q3 v3
```

- Scan the table: scan

```
row f1:q1 [] v1
row f2:q2 [] v2
row f3:q3 [] v3
```

- Delete the test table: root@accumulo test> deletetable
Table: [test] has been deleted.

Example

The following example shows the verification of an unsecured OpDB installation:

```
root@accumulo> createtable test
root@accumulo test> userpermissions
System permissions: System.GRANT, System.CREATE_TABLE, System.DROP_TABLE, System.ALTER_TABLE, System.CREATE_USER, System.DROP_USER, System.ALTER_USER, System.SYSTEM, System.CREATE_NAMESPACE, System.DROP_NAMESPACE, System.ALTER_NAMESPACE, System.OBTAIN_DELEGATION_TOKEN

Namespace permissions (accumulo): Namespace.READ, Namespace.ALTER_TABLE

Table permissions (accumulo.metadata): Table.READ, Table.ALTER_TABLE
Table permissions (accumulo.replication): Table.READ
Table permissions (accumulo.root): Table.READ, Table.ALTER_TABLE
Table permissions (test): Table.READ, Table.WRITE, Table.BULK_IMPORT, Table.ALTER_TABLE, Table.GRANT, Table.DROP_TABLE, Table.GET_SUMMARIES
```

```
root@accumulo test> insert row f1 q1 v1
root@accumulo test> insert row f2 q2 v2
root@accumulo test> insert row f3 q3 v3
root@accumulo test> scan
row f1:q1 [] v1
row f2:q2 [] v2
row f3:q3 [] v3
root@accumulo test> deletetable
deletetable { test } (yes|no)? yes
Table: [test] has been deleted.
root@accumulo>
```

Installing Accumulo Parcel 1.1.0

Learn how to install the Operational Database powered by Apache Accumulo (Accumulo-on-cdp 1.1.0 parcel) on your CDP Private Cloud Base environment.

Install Operational Database powered by Apache Accumulo

Before installing Operational Database powered by Apache Accumulo (OpDB), you first have to install Cloudera Manager, CDP cluster services, then install the OpDB CSD and the OpDB Parcel. Ensure you verify your installation by running certain basic commands.

Install Operational Database powered by Apache Accumulo CSD file

As Operational Database powered by Apache Accumulo (OpDB) is an add-on service, you have to install its Custom Services Descriptor (CSD) file to manage the Accumulo on CDP service.

About this task

A Custom Services Descriptor (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 files, see the *Custom Service Descriptor Files* section in the *Add-on Services* documentation.

Before you begin

Install Cloudera Manager.

Procedure

1. Download the ACCUMULO_CDP-1.1.0 jar file for the CSD that provides the Accumulo on CDP service.

You can download that JAR file from:

```
archive.cloudera.com/p/accumulo7/1.1.0/csd
```

2. Copy the ACCUMULO_CDP-1.1.0.jar file to the location where Cloudera Manager is configured to look for the CSD files. That is by default in the `/opt/cloudera/csd` folder on the host where Cloudera Manager is running.
3. Restart Cloudera Manager as described in the *Configuring the Location of Customer Service* section in the *Add-on Services* documentation.

What to do next

Install CDP cluster and services.

Related Information

[Add-on Services](#)

Install CDP

Before you install the Operational Database powered by Apache Accumulo (OpDB) parcel, you have to install CDP Private Cloud Base and add the mandatory services.

The OpDB service is only supported on CDP Private Cloud Base. For full cluster installation Cloudera recommends following the [CDP Private Cloud Base Installation Guide](#).

For a FIPS compliant OpDB service you have to follow the FIPS deployment CDP setup steps. For more information, see [Installing and Configuring CDP with FIPS](#).

Ensure you install the following mandatory services for OpDB powered by Apache Accumulo:

- HDFS
- ZooKeeper

Install Operational Database powered by Apache Accumulo parcel

After you have installed the Operational Database powered by Apache Accumulo (OpDB) CSD and the required CDP cluster and services, you have to install the OpDB parcel.

There are two options to install the OpDB parcel:

- Using a Local Parcel Repository
- Using Remote Parcel Repository URLs (archive.cloudera.com)

Install Operational Database powered by Apache Accumulo parcel using Local Parcel Repository

You can install the OpDB powered by Apache Accumulo Parcel (OpDB) using your Local Parcel Repository. For this, you have to download the parcel, copy it to your Local Parcel Repository, and then distribute and activate it using Cloudera Manager.

Before you begin

- Install the OpDB CSD file
- Install CDP and add the HDFS and ZooKeeper services to your deployment.

Procedure

1. Download the ACCUMULO_ON_CDP-1.1.0.p0.11738502-el7.parcel OpDB parcel from:

```
archive.cloudera.com/p/accumulo7/1.1.0/parcels
```

2. Copy the ACCUMULO_ON_CDP-1.1.0.p0.11738502-el7.parcel file to the Local Parcel Repository. The Local Parcel Repository is by default located in `/opt/cloudera/parcel-repo` directory on the host where Cloudera Manager is running.

- In Cloudera Manager, navigate to Parcels.

The screenshot shows the Cloudera Manager interface. On the left, a dark sidebar contains a navigation menu with items: Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, Private Cloud (with a 'New' badge), Parcels (highlighted with an orange arrow), Running Commands, and Support. The main content area is titled 'Home' and shows a navigation bar with 'Status', 'All Health Issues', 'Configuration', and 'All Recent Commands'. Below this, a section for 'Cluster 1' is displayed, showing a list of installed parcels. Each parcel row includes a green checkmark, an icon, and the parcel name. The parcels listed are: 4 Hosts, ATLAS-1, HBASE-1, HDFS-1, HIVE-1, HIVE_ON_TEZ-1, HUE-1, IMPALA-1, KAFKA-1, KNOX-1, KUDU-1, LIVY-1, OOZIE-1, and OZONE-1.

- Click Check for New Parcels.
The ACCUMULO_ON_CDP-1.1.0.p0.11738502-el7.parcel parcel appears in the Parcel list.
- In the row of ACCUMULO_ON_CDP-1.1.0.p0.11738502-el7.parcel, click Distribute.
- After the distribution is completed, click Activate.

Results

OpDB parcel is installed.

What to do next

Add the Accumulo on CDP service.

Install Operational Database powered by Apache Accumulo using Remote Parcel Repository

You can install the Operational Database powered by Apache Accumulo (OpDB) parcel using the archive.cloudera.com Remote Parcel Repository.

Before you begin


- Install the OpDB CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.

Procedure

1. In Cloudera Manager, navigate to Parcels.

The screenshot shows the Cloudera Manager interface. On the left is a dark sidebar with the Cloudera Manager logo and a search bar. Below the search bar are several menu items: Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, Private Cloud (with a 'New' badge), Parcels (highlighted with an orange arrow), Running Commands, and Support. The main content area is titled 'Home' and has a navigation bar with 'Status', 'All Health Issues', 'Configuration' (with a dropdown arrow), and 'All Recent Comm'. Below the navigation bar, there is a section for 'Cluster 1' (indicated by a green checkmark) showing 'Cloudera Runtime 7.1.5 (Parcels)'. A list of services follows, each with a green checkmark and a vertical ellipsis menu icon: 4 Hosts, ATLAS-1, HBASE-1, HDFS-1, HIVE-1, HIVE_ON_TEZ-1, HUE-1, IMPALA-1, KAFKA-1, KNOX-1, KUDU-1, LIVY-1, OOZIE-1, and OZONE-1.

2. Click Parcel Repositories & Network Settings.



3. In the Remote Parcel Repository URLs section, click (Add Another) to include another remote parcel repository URL.

4. Add the following URL:

```
https://archive.cloudera.com/p/accumulo7/1.1.0/parcels/
```

5. Click Save & Verify Configuration.

6. In the row of ACCUMULO_ON_CDP 1.1.0.p0.11738502-el7.parcel, click Download.

7. After the download is completed, click Distribute.

8. Once the distribution is completed, click Activate.

What to do next

Add the Accumulo on CDP service.

Add Accumulo on CDP service

After you have installed the Operational Database powered by Apache Accumulo (OpDB) CSD, the required CDP cluster and services, and the OpDB parcel, you can add the Accumulo on CDP service to your cluster.

You can add a secure or an unsecured Accumulo on CDP service to your CDP Private Cloud Base deployment when installing OpDB. A secure Accumulo on CDP service requires a secure cluster, which means an SASL-enabled Kerberized cluster.

The steps to add a secure Accumulo on CDP service are different from the steps to add an unsecure Accumulo on CDP service. If you require a FIPS compliant OpDB service, follow the steps in the Add secure Accumulo on CDP service to your cluster.

Add unsecure Accumulo on CDP service to your cluster

If your cluster is not an SASL-enabled Kerberized cluster, add an unsecure Accumulo on CDP service to your cluster when installing an Operational Database powered by Apache Accumulo (OpDB).


Before you begin

- Install the OpDB CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.
- Install the OpDB parcel.

Procedure

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

- Click the action menu of the cluster to which you want to add the Accumulo on CDP service (for example, Cluster 1), and select Add Service.



- Select Accumulo on CDP and click Continue.
- Select the dependent services and click Continue.
- Assign the Accumulo roles to the hosts in your cluster:
 - Assign a Table Server role on each host to which the DataNode role is assigned.
 - Assign the Monitor, Garbage Collector, and Master role to non-DataNodes.
 - Assign the Gateway role to any hosts where you want to use the Accumulo on CDP service, and where you do not already have any other OpDB roles assigned.
- Click Continue.
- Configure the Accumulo Instance Secret.



Important: Cloudera recommends to protect this information, because the security of the Accumulo service relies on it.

- Set the Accumulo Instance Name.
- If this is the first installation of OpDB, ensure that the Initial Accumulo option is selected.
- Click Continue.
- Wait while Cloudera Manager does the initial service setup.
- Click Continue.
- Click Finish.

Add secure Accumulo on CDP service to your cluster

To ensure data availability to diverse sets of users with varying levels of permissions and security clearance, add a secure Accumulo on CDP service.

About this task

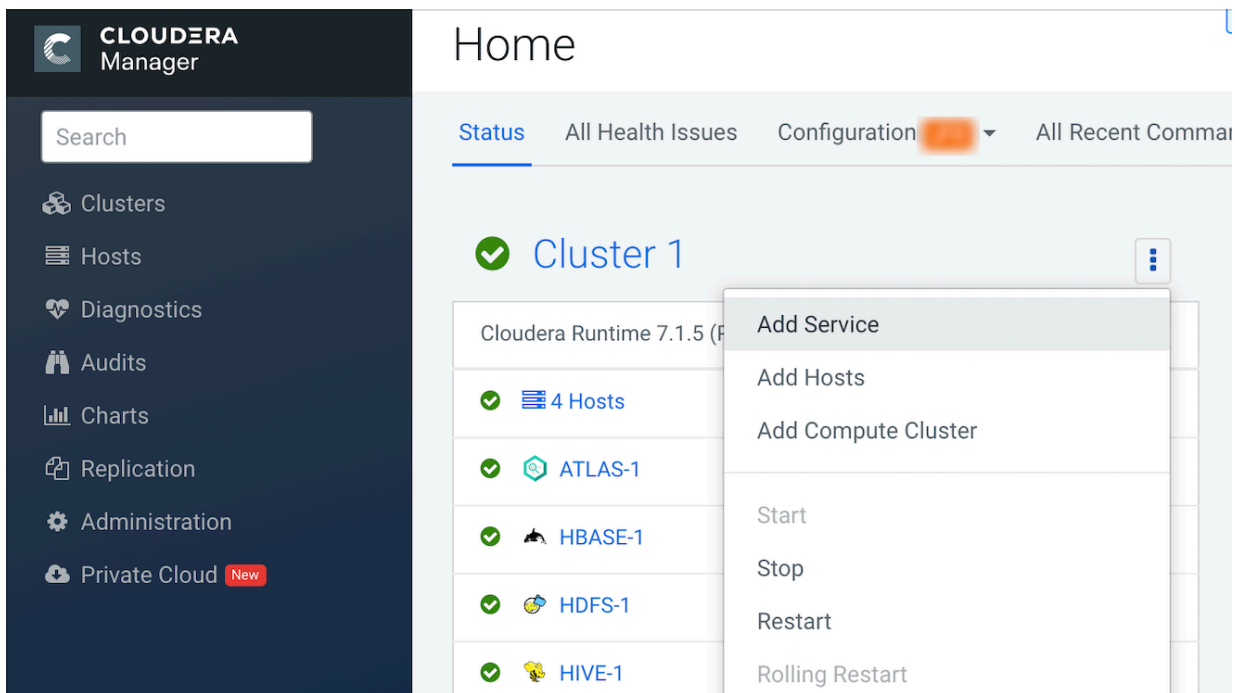
If your secure Accumulo on CDP service needs to be FIPS compliant, ensure that you have followed the FIPS specific deployment setup steps provided in [Installing and Configuring CDP with FIPS](#).

Before you begin

- Install OpDB CSD file.
- Install CDP.
- Add the HDFS and ZooKeeper services to your CDP deployment.
- Install the OpDB parcel.
- Ensure that you have a SASL-enabled Kerberized cluster.

Procedure

1. Navigate to the Cloudera Manager Admin Console home page.
2. Select the action menu of the cluster to which you want to add the Accumulo on CDP service (for example, Cluster 1), and select Add Service.



3. Select Accumulo on CDP and click Continue.
4. Select the dependent services and click Continue.
5. Assign the Accumulo roles to the hosts in your cluster:
 - Assign a Table Server role on each host to which the DataNode role is assigned.
 - Assign the Monitor, Garbage Collector, and Master role to non-DataNodes.
 - Assign the Gateway role to any hosts here you want to use the OpDB service and that do not already have other OpDB roles assigned.
6. Click Continue.
7. Configure the Accumulo Instance Secret.



Important: Cloudera recommends to protect this information, because the security of the Accumulo service relies on it.

8. Set the Accumulo Instance Name.
9. Select Enable Kerberos Authentication.
10. If this is the first installation of OpDB, ensure that the Initial Accumulo option is selected.
11. Click Continue.
12. Wait while the Cloudera Manager does the initial service setup.

13. Click Continue.
14. Click Finish.
15. Create a client configuration file for the admin user.

This step needs to be performed on each Gateway machine where you will run an OpDB client.

Client configuration can be created by editing the configuration file in the following way:

```
auth.principal=accumulo/[***HOSTNAME***]
auth.type=kerberos
auth.token=/var/run/cloudera-scm-agent/process/[***TIMESTAMP***]-[***SERVICE NAME***]-[***INSTANCE NAME***]/accumulo_on_cdp.keytab
sasl.enabled=true
sasl.qop=auth-conf
```

By default the Accumulo principal on which master is running becomes the administrator principal. You can set up a different administrative user. For more information, see the [Administrative User section of the user manual](#). In this case, you need to change the `auth.principal` and `auth.token` configuration properties accordingly.

Verify your Operational Database powered by Apache Accumulo installation

You can verify your Operational Database powered by Apache Accumulo (OpDB) installation using the Accumulo shell.

Before you begin

- Install the OpDB CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.
- Install the OpDB parcel.
- Add the Accumulo on CDP service to your deployment.

Procedure

1. Launch the Accumulo shell for the admin user.

- Unsecure OpDB

```
$ accumulo shell -u root
Password: *****
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
root@accumulo>
```

- Secure OpDB

```
kinit -kt /var/run/cloudera-scm-agent/process/[***TIMESTAMP***]-[***SERVICE NAME***]-[***INSTANCE NAME***]/accumulo_on_cdp.keytab accumulo/${***HOSTNAME***}
accumulo shell
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
```

```
- type 'help' for a list of available commands
-
accumulo/hostname@REALM.COM@accumulo>
```

2. Run the following commands to verify your installation:

a) Create a table: createtable test

The test table is created and becomes the current table.

b) List the admin user's permission and ensure it has admin and system level privileges: userpermissions

c) On a secure cluster, the user who creates the table does not get access to the new table by default. These permissions have to be set manually:

```
'grant Table.READ -t test -u accumulo/hostname@REALM.COM'
'grant Table.WRITE -t test -u accumulo/hostname@REALM.COM'
'grant Table.ALTER_TABLE -t test -u accumulo/hostname@REALM.COM'
'grant Table.DROP_TABLE -t test -u accumulo/hostname@REALM.COM'
```

d) Insert some rows to the test table:

```
insert row f1 q1 v1
insert row f2 q2 v2
insert row f3 q3 v3
```

e) Scan the table: scan

```
row f1:q1 [] v1
row f2:q2 [] v2
row f3:q3 [] v3
```

f) Delete the test table: root@accumulo test> deletetable

Table: [test] has been deleted.

Example

The following example shows the verification of an unsecured OpDB installation:

```
root@accumulo> createtable test
root@accumulo test> userpermissions
System permissions: System.GRANT, System.CREATE_TABLE, System.DROP_TABLE, System.ALTER_TABLE, System.CREATE_USER, System.DROP_USER, System.ALTER_USER, System.SYSTEM, System.CREATE_NAMESPACE, System.DROP_NAMESPACE, System.ALTER_NAMESPACE, System.OBTAIN_DELEGATION_TOKEN

Namespace permissions (accumulo): Namespace.READ, Namespace.ALTER_TABLE

Table permissions (accumulo.metadata): Table.READ, Table.ALTER_TABLE
Table permissions (accumulo.replication): Table.READ
Table permissions (accumulo.root): Table.READ, Table.ALTER_TABLE
Table permissions (test): Table.READ, Table.WRITE, Table.BULK_IMPORT, Table.ALTER_TABLE, Table.GRANT, Table.DROP_TABLE, Table.GET_SUMMARIES
root@accumulo test> insert row f1 q1 v1
root@accumulo test> insert row f2 q2 v2
root@accumulo test> insert row f3 q3 v3
root@accumulo test> scan
row f1:q1 [] v1
row f2:q2 [] v2
row f3:q3 [] v3
root@accumulo test> deletetable
deletetable { test } (yes|no)? yes
Table: [test] has been deleted.
root@accumulo>
```

Installing Accumulo Parcel 1.10

Learn how to install the Accumulo parcel on your CDP Private Cloud Base environment.

Install Accumulo

Before installing the Accumulo, you first have to install Cloudera Manager and CDP cluster services, after that install the Accumulo CSD and the Accumulo Parcel. Ensure you verify your installation by running certain basic commands.

Install Accumulo CSD file

As Accumulo is an add-on service, you have to install its Custom Services Descriptor (CSD) file to manage the Accumulo on CDP service.

About this task

A Custom Services Descriptor (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 files, see the *Custom Service Descriptor Files* section in the *Add-on Services* documentation.

Before you begin

Install Cloudera Manager.

Procedure

1. Download the ACCUMULO_CDP-1.10.3_7.1.7.2000.0 jar file for the CSD that provides the Accumulo on CDP service.
You can download that JAR file from:
https://archive.cloudera.com/p/accumulo7/1.10.3_7.1.7.2000.0/
2. Copy the ACCUMULO_CDP-1.10.3_7.1.7.2000.0 file to the location where Cloudera Manager is configured to look for the CSD files. That is by default in the /opt/cloudera/csd folder on the host where Cloudera Manager is running.
3. Restart Cloudera Manager as described in the *Configuring the Location of Customer Service* section in the *Add-on Services* documentation.

What to do next

Install CDP cluster and services.

Related Information

[Add-on Services](#)

Install CDP

Before you install the Accumulo parcel, you have to install CDP Private Cloud Base and add the mandatory services.

The Accumulo service is only supported on CDP Private Cloud Base. For full cluster installation Cloudera recommends following the [CDP Private Cloud Base Installation Guide](#).

For a FIPS-compliant Accumulo service, you have to follow the FIPS deployment CDP setup steps. For more information, see [Installing and Configuring CDP with FIPS](#).

Ensure that you install the following mandatory services for Accumulo:

- HDFS
- ZooKeeper

Install Accumulo 1.10 parcel

After you have installed the Accumulo CSD and the required CDP cluster and services, you have to install the Accumulo parcel.

There are two options to install the Accumulo parcel:

- Using a Local Parcel Repository
- Using Remote Parcel Repository URLs (archive.cloudera.com)

Install Accumulo parcel using Local Parcel Repository

You can install the Accumulo Parcel using your Local Parcel Repository. For this, you have to download the parcel, copy it to your Local Parcel Repository, and then distribute and activate it using Cloudera Manager.

Before you begin

- Install the Accumulo CSD file
- Install CDP and add the HDFS and ZooKeeper services to your deployment.

Procedure

1. Download the ACCUMULO_ON_CDP-1.10.3_7.1.7.2000.0.p0.43785444-el8.parcel Accumulo parcel from:
https://archive.cloudera.com/p/accumulo7/1.10.3_7.1.7.2000.0/
2. Copy the ACCUMULO_ON_CDP-1.10.3_7.1.7.2000.0.p0.43785444-el8.parcel file to the Local Parcel Repository. The Local Parcel Repository is by default located in `/opt/cloudera/parcel-repo` directory on the host where Cloudera Manager is running.

- In Cloudera Manager, navigate to Parcels.

The screenshot shows the Cloudera Manager interface. On the left is a dark sidebar with the Cloudera Manager logo and a search bar. Below the search bar are navigation links: Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, and Private Cloud (with a 'New' badge). At the bottom of the sidebar are links for Parcels, Running Commands, and Support. An orange arrow points to the 'Parcels' link. The main content area is titled 'Home' and has tabs for Status, All Health Issues, Configuration, and All Recent Commands. The 'Status' tab is active, showing a green checkmark and 'Cluster 1'. Below this is a list of installed parcels for Cloudera Runtime 7.1.5 (Parcels). Each row includes a green checkmark, an icon, and the parcel name. The parcels listed are: 4 Hosts, ATLAS-1, HBASE-1, HDFS-1, HIVE-1, HIVE_ON_TEZ-1, HUE-1, IMPALA-1, KAFKA-1, KNOX-1, KUDU-1, LIVY-1, OOZIE-1, and OZONE-1.

- Click Check for New Parcels.
The `ACCUMULO_ON_CDP-1.10.3_7.1.7.2000.0.p0.43785444-e18.parcel` parcel appears in the Parcel list.
- In the row of `ACCUMULO_ON_CDP-1.10.3_7.1.7.2000.0.p0.43785444-e18.parcel`, click Distribute.
- After the distribution is completed, click Activate.

Results

The Accumulo parcel is installed.

What to do next

Add the Accumulo on CDP service.

Install Accumulo using Remote Parcel Repository

You can install the Accumulo parcel using the archive.cloudera.com Remote Parcel Repository.

Before you begin

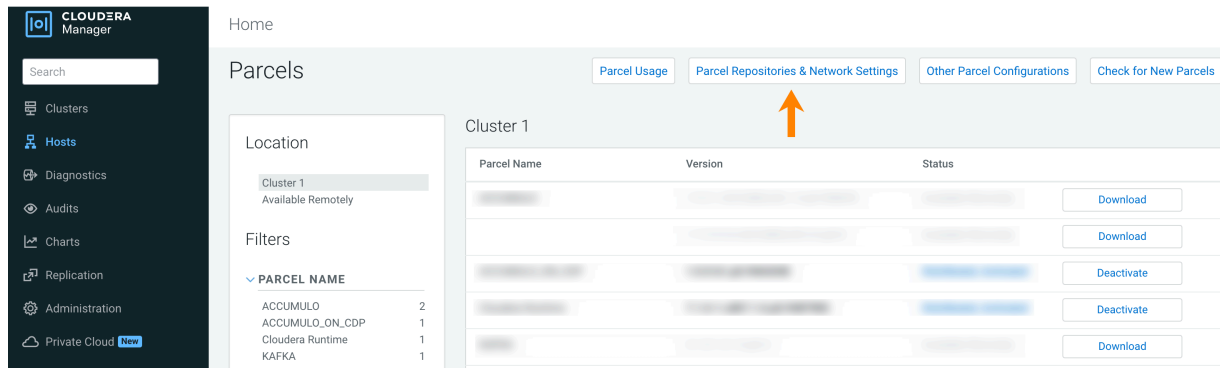
- Install the Accumulo CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.

Procedure

1. In Cloudera Manager, navigate to Parcels.

The screenshot shows the Cloudera Manager interface. On the left is a dark sidebar with the Cloudera Manager logo and a search bar. Below the search bar are several menu items: Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, and Private Cloud (with a 'New' badge). At the bottom of the sidebar, there are icons for Parcels, Running Commands, and Support. An orange arrow points to the 'Parcels' menu item. The main content area is titled 'Home' and has a navigation bar with 'Status', 'All Health Issues', 'Configuration', and 'All Recent Comm'. Below the navigation bar, there is a section for 'Cluster 1' with a green checkmark icon. Underneath, it shows 'Cloudera Runtime 7.1.5 (Parcels)' and a list of services, each with a green checkmark and a vertical ellipsis menu icon. The services listed are: 4 Hosts, ATLAS-1, HBASE-1, HDFS-1, HIVE-1, HIVE_ON_TEZ-1, HUE-1, IMPALA-1, KAFKA-1, KNOX-1, KUDU-1, LIVY-1, OOZIE-1, and OZONE-1.

2. Click Parcel Repositories & Network Settings.



3. In the Remote Parcel Repository URLs section, click **+** (Add Another) to include another remote parcel repository URL.
4. Add the following URL:

https://archive.cloudera.com/p/accumulo7/1.10.3_7.1.7.2000.0/
5. Click Save & Verify Configuration.
6. In the row of ACCUMULO_ON_CDP 1.10.0.p0.11738502-el7.parcel, click Download.
7. After the download is completed, click Distribute.
8. Once the distribution is completed, click Activate.

What to do next

Add the Accumulo on CDP service.

Add Accumulo on CDP service

After you have installed the Accumulo CSD, the required CDP cluster and services, and the Accumulo parcel, you can add the Accumulo on CDP service to your cluster.

You can add a secure or an unsecured Accumulo on CDP service to your CDP Private Cloud Base deployment when installing Accumulo. A secure Accumulo on CDP service requires a secure cluster, which means an SASL-enabled Kerberized cluster.

The steps to add a secure Accumulo on CDP service are different from the steps to add an unsecure Accumulo on CDP service. If you require a FIPS compliant Accumulo service, follow the steps in the Add secure Accumulo on CDP service to your cluster.

Add unsecure Accumulo on CDP service to your cluster

If your cluster is not an SASL enabled Kerberized cluster, add an unsecure Accumulo on CDP service to your cluster when installing Accumulo.


Before you begin

- Install Accumulo CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.
- Install the Accumulo parcel.

Procedure

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

- Click the action menu of the cluster to which you want to add the Accumulo on CDP service (for example, Cluster 1), and select Add Service.



- Select Accumulo on CDP and click Continue.
- Select the dependent services and click Continue.
- Assign the Accumulo roles to the hosts in your cluster:
 - Assign a Table Server role on each host to which the DataNode role is assigned.
 - Assign the Monitor, Garbage Collector, Tracer, and Master role to non-DataNodes.
 - Assign the Gateway role to any hosts where you want to use the Accumulo on CDP service, and where you do not already have any other Accumulo roles assigned.
- Click Continue.
- Configure the Accumulo Instance Secret.



Important: Cloudera recommends to protect this information, because the security of the Accumulo service relies on it.

- Set the Accumulo Instance Name.
- Set the Trace User and Trace Password settings.

It is recommended that you not leave the Trace User set to root, because this is not a secure configuration.
- If this is the first installation of Accumulo, ensure that the Initial Accumulo option is selected.
- Click Continue.
- Wait while the Cloudera Manager does the initial service setup.
- Click Continue.
- Click Finish.

What to do next

Ensure that the trace user and trace table exist.

Add secure Accumulo on CDP service to your cluster

To ensure data availability to diverse sets of users with varying levels of permissions and security clearance, add a secure Accumulo on CDP service.

About this task


If your secure Accumulo on CDP service needs to be FIPS compliant, ensure that you have followed the FIPS specific deployment setup steps provided in [Installing and Configuring CDP with FIPS](#).

Before you begin

- Install Accumulo CSD file.
- Install CDP.
- Add the HDFS and ZooKeeper services to your CDP deployment.
- Install the Accumulo parcel.
- Ensure that you have a SASL enabled Kerberized cluster.

Procedure

1. Navigate to the Cloudera Manager Admin Console home page.
2. Select the action menu of the cluster to which you want to add the Accumulo on CDP service (for example, Cluster 1), and select Add Service.



3. Select Accumulo on CDP and click Continue.
4. Select the dependent services and click Continue.
5. Assign the Accumulo roles to the hosts in your cluster:
 - Assign a Table Server role on each host to which the DataNode role is assigned.
 - Assign the Monitor, Garbage Collector, Tracer, and Master role to non-DataNodes. Ensure that the Tracer Role is on the same node as the Master.
 - Assign the Gateway role to any hosts here you want to use the Accumulo service and that do not already have other Accumulo roles assigned.
6. Click Continue.
7. Configure the Accumulo Instance Secret.



Important: Cloudera recommends to protect this information, because the security of the Accumulo service relies on it.

8. Set the Accumulo Instance Name.

9. Select Enable Kerberos Authentication.



Note: This setting overwrites the trace user with the default Accumulo Kerberos principal. This is the reason the Tracer Role must be on the same node as the Master.

10. If this is the first installation of Accumulo, ensure that the Initial Accumulo option is selected.

11. Click Continue.

12. Wait while Cloudera Manager does the initial service setup.

13. Click Continue.

14. Click Finish.

15. Restart all role except Tracers.

16. Create a client configuration file for the admin user.

This step needs to be performed on each Gateway machine where you will run an Accumulo client.

Client configuration can be created by editing the configuration file in the following way:

```
instance.rpc.sasl.enabled=true
rpc.sasl.qop=auth-conf
kerberos.server.primary=accumulo
```

By default, the Accumulo principal on which master is running becomes the administrator principal. You can set up a different administrative user. For more information, see the [Administrative User section of the user manual](#). In this case you need to change the `auth.principal` and `auth.token` configuration properties accordingly.

What to do next

Ensure that the trace table exists.

Customizing Kerberos Principals

The custom Kerberos principals is an advanced optional feature that allows you to customize the Kerberos principals for the Accumulo service. Learn how to customize the Kerberos principals in the Cloudera Manager.

About this task

- Cloudera Manager configures CDP services to use the default Kerberos principal names. Cloudera recommends not changing the default Kerberos principal. Do not change the Kerberos principal for existing clusters and services.
- If your security policies require you to customize the Kerberos principals for services, contact Cloudera Professional Services for additional help. Configuring the recommended solution requires extra steps and automation outside CDP to succeed.
- Use the following steps only when adding a new service to a new cluster created by enabling custom Kerberos principals using Cloudera Manager. Customizing a system username is not supported at the moment.

Before you begin

- Ensure that you are using the latest version of Cloudera Manager. The minimum supported version is CM 7.11.3.0.
- Before customizing the Kerberos principals in the Cloudera Manager, enable the feature flag for custom Kerberos principals before setting up the cluster. See *Enabling feature flag for Custom Kerberos Principals and System Users*.
- Before adding a new service with custom Kerberos principal, you must use the Ranger UI and update the appropriate policies.

- Add the following auth-to-local rule in the Additional Rules to Map Kerberos Principals to Short Names property in the HDFS service using Cloudera Manager.

```
RULE:[2:$1/$2@$0](<custom_accumulo_princiapl>/.*@<realm_name>)s/./*/accumulo/
```

Replace the `<custom_accumulo_princiapl>` and `<realm_name>` options with the actual values you plan to use during the Accumulo installation.

For more information on adding the auth-to-local rule, see *Using auth-to-local rules to isolate cluster users*.

Procedure

1. Navigate to Custom Kerberos Principals and System Users after the Review Changes operation, for customizing the Kerberos principals.

The screenshot shows the Cloudera Manager interface for configuring 'Add Accumulo on CDP Service to Cluster 1'. The left sidebar shows the navigation menu with 'Custom Kerberos Principals and System Users' selected. The main content area displays the configuration page with a warning message and a form to customize Kerberos principals and system users.

Custom Kerberos Principals and System Users

Warning: This is an advanced feature that allows you to customize the kerberos principals and the system users that the services will run as. Note: Default system users will be created automatically, but all non default users must be first created manually on all the hosts. All the necessary files/directories will have permissions set for these users. These permissions are very difficult to change afterwards. If you are unsure what to do, leave them unmodified.

Customize Kerberos Principals
 Customize System Users

Kerberos Principal: Accumulo on CDP (Service-Wide) [Undo](#)
 1 - 1 of 1

[Generate Dependent Configurations](#)

Clicking this button will generate custom rules to map service principals and users that are not matching. Cloudera recommends that service principals and users match when being set. Following this recommendation will not generate any rules in the box.

Navigation buttons: [Cancel](#), [Back](#), [Continue](#)

- Select the Customize Kerberos Principals option as seen in the following figure.

- Update the Kerberos principal for services.

You can change the names as required. You must use the same name for a Kerberos principal already set in the auth-to-local configuration. For example, use *accumulo-2* as a principal value.

- Click Generate Dependent Configurations.



Note: The resulting rule must match the rule added to the auth-to-local configuration mentioned under the *Before you begin* section.

- Click Continue.

Creating a trace user in unsecure Accumulo deployment

After you add an unsecure Accumulo on CDP service, you have to ensure that the trace user exists. This step is only required if you have an unsecure Accumulo deployment.

About this task

This step is not required if you have a secure Accumulo deployment, because in that case the trace user is automatically created when the Accumulo on CDP service is added to the cluster.

Before you begin

- Install the Accumulo CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.
- Install the Accumulo parcel.
- Add an unsecure Accumulo on CDP to your deployment.

Ensure that you have added an unsecure Accumulo service to your cluster.

Procedure

1. Launch the Accumulo shell for the admin user.

```
$ accumulo shell -u root
Password: *****
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
root@accumulo>
```

2. List the existing users and check if trace user exists: users
If the trace user exists, you can continue your Accumulo installation.
3. If the trace user does not exist, create it with the following command: createuser trace

What to do next

Ensure that the trace table exists.

Check trace table

After you have added the Accumulo on CDP service as a part of the Accumulo installation, you have to check whether the trace table, that records the time span of operations, exists.

About this task

Accumulo records the time that various operations take when tracing is turned on. It follows all the requests made on behalf of the user throughout the distributed infrastructure of Accumulo, and across all threads of execution.

These timespans are inserted into the trace table in Accumulo. You can browse recent traces from the Accumulo monitor page. You can also read the trace table directly like any other table.

Before you begin

- Install the Accumulo CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.
- Install the Accumulo parcel.
- Add the Accumulo on CDP service to your deployment.

Procedure

1. Launch the Accumulo shell for the admin user.

- Unsecure Accumulo

```
$ accumulo shell -u root
Password: *****
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
root@accumulo>
```

- Secure Accumulo

```
kinit -kt /var/run/cloudera-scm-agent/process/[***TIMESTAMP***]-[***SERVICE NAME***]-[***INSTANCE NAME***]/accumulo_on_cdp.keytab accumulo/[***HOSTNAME***]
accumulo shell
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
accumulo/hostname@REALM.COM@accumulo>
```

2. List tables and check if the trace table exists.

```
> tables
accumulo.metadata
accumulo.replication
accumulo.root
trace
>
```

3. If the trace table does not exist, create it with the following command: createtable trace.

What to do next

Provide user permissions.

Provide user permissions

After the Accumulo on CDP service is added as part of the Accumulo, you have to provide privileges to the principal in which the Tracer is running.

Before you begin

- Install the Accumulo CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.
- Install the Accumulo parcel.
- Add the Accumulo on CDP service.
- Check the trace table.

Procedure

1. Launch the Accumulo shell for the admin user.

- Unsecure Accumulo

```
$ accumulo shell -u root
Password: *****
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
root@accumulo>
```

- Secure Accumulo

```
kinit -kt /var/run/cloudera-scm-agent/process/[***TIMESTAMP***]-[***SERVICE NAME***]-[***INSTANCE NAME***]/accumulo_on_cdp.keytab accumulo/[***HOSTNAME***]
accumulo shell
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
accumulo/hostname@REALM.COM@accumulo>
```

2. Add READ, WRITE, and ALTER_TABLE privileges to the principal in which the Tracer is running.

- Unsecure Accumulo

```
grant Table.READ -t trace -u trace
grant Table.WRITE -t trace -u trace
grant Table.ALTER_TABLE -t trace -u trace
```

- Secure Accumulo

```
'grant Table.READ -t trace -u accumulo/hostname@REALM.COM'
'grant Table.WRITE -t trace -u accumulo/hostname@REALM.COM'
'grant Table.ALTER_TABLE -t trace -u accumulo/hostname@REALM.COM'
```

In an unsecure Accumulo deployment if you have configured the trace user with a `_HOST` instance component, you must repeat this command with an instance for each host that will run a Tracer or Monitor role.

3. Start all Tracer roles.

What to do next

Verify your Accumulo installation.

Verify your Accumulo installation

You can verify your Accumulo installation using the Accumulo shell.

Before you begin

- Install the Accumulo CSD file.
- Install CDP and add the HDFS and ZooKeeper services to your deployment.
- Install the Accumulo parcel.
- Add the Accumulo on CDP service to your deployment.
- Check trace table.
- Provide user permissions.

Procedure

1. Launch the Accumulo shell for the admin user.

- Unsecure Accumulo

```
$ accumulo shell -u root
Password: *****
Shell - Apache Accumulo Interactive Shell
-
```



```

- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
root@accumulo>

```

- Secure Accumulo

```

kinit -kt /var/run/cloudera-scm-agent/process/[***TIMESTAMP***]-[***SERVICE NAME***]-[***INSTANCE NAME***]/accumulo_on_cdp.keytab accumulo/${***HOSTNAME***}
accumulo shell
Shell - Apache Accumulo Interactive Shell
-
- version: [***ACCUMULO VERSION NUMBER***]
- instance name: [***ACCUMULO INSTANCE NAME***]
- instance id: [***ACCUMULO INSTANCE ID***]
-
- type 'help' for a list of available commands
-
accumulo/hostname@REALM.COM@accumulo>

```

2. Run the following commands to verify your installation:

- Create a table: `createtable test`
The test table is created and becomes the current table.
- Grant permission to READ or WRITE on the newly created table. For more information, see *Provide user permissions*.
- Run the following command to provide an additional permission to delete a table.

```
grant Table.DROP_TABLE -t test -u accumulo/hostname@REALM.COM
```

- List the admin user's permission and ensure it has admin and system level privileges: `userpermission`
- Insert some rows to the test table:

```
insert row f1 q1 v1
insert row f2 q2 v2
insert row f3 q3 v3
```

- Scan the table: `scan`

```
row f1:q1 [] v1
row f2:q2 [] v2
row f3:q3 [] v3
```

- Delete the test table: `root@accumulo test> deletetable`
Table: [test] has been deleted.

Example

The following example shows the verification of an unsecured Accumulo installation:

```

root@accumulo> createtable test
root@accumulo test> userpermissions
System permissions: System.GRANT, System.CREATE_TABLE, System.DROP_TABLE, System.ALTER_TABLE, System.CREATE_USER, System.DROP_USER, System.ALTER_USER, System.SYSTEM, System.CREATE_NAMESPACE, System.DROP_NAMESPACE, System.ALTER_NAMESPACE, System.OBTAIN_DELEGATION_TOKEN

Namespace permissions (accumulo): Namespace.READ, Namespace.ALTER_TABLE

Table permissions (accumulo.metadata): Table.READ, Table.ALTER_TABLE

```

```

Table permissions (accumulo.replication): Table.READ
Table permissions (accumulo.root): Table.READ, Table.ALTER_TABLE
Table permissions (test): Table.READ, Table.WRITE, Table.BULK_IMPORT, Table.ALTER_TABLE, Table.GRANT, Table.DROP_TABLE, Table.GET_SUMMARIES
Table permissions (trace): Table.READ, Table.WRITE, Table.BULK_IMPORT, Table.ALTER_TABLE, Table.GRANT, Table.DROP_TABLE, Table.GET_SUMMARIES
root@accumulo test> insert row f1 q1 v1
root@accumulo test> insert row f2 q2 v2
root@accumulo test> insert row f3 q3 v3
root@accumulo test> scan
row f1:q1 [] v1
row f2:q2 [] v2
row f3:q3 [] v3
root@accumulo test> deletetable
deletetable { test } (yes|no)? yes
Table: [test] has been deleted.
root@accumulo>

```

Upgrading Accumulo from 1.10.3 to 2.1.2

You can upgrade from Accumulo version 1.10.3 to Accumulo 2.1.2. Additionally, you must remove and reinstall the Accumulo service as a part of upgrading.

Upgrading CDP Private Cloud Base

You must upgrade your cluster from CDP Private Cloud Base 7.1.7.SP2 to CDP Private Cloud Base 7.1.9 SP1. For information on upgrading the CDP Private Cloud Base cluster, see [Upgrading CDP Private Cloud Base to a higher version](#).

Removing Accumulo 1.10.3

You must remove Accumulo 1.10.3 installed on your CDP cluster before upgrading the Accumulo service.

Before you begin

- You must stop all the ongoing administrative operations before you remove the Accumulo service.
- Make a note of all the non-default configurations before starting the upgrade process.

Procedure

1. Run the `root@accumulo> fate print` command to check for any Fault Tolerant Executor (FATE) transactions.
2. Run the `root@accumulo> fate delete <transaction id>` command to remove such transactions. Ensure tablet servers and the master are healthy.
3. Log in to Cloudera Manager Clusters Accumulo
4. Click Actions Stop .
5. Click Configuration.
6. Check all the parameters under the Configuration tab and identify and note down any non-default settings. This information is required when you add the Accumulo service to your cluster.



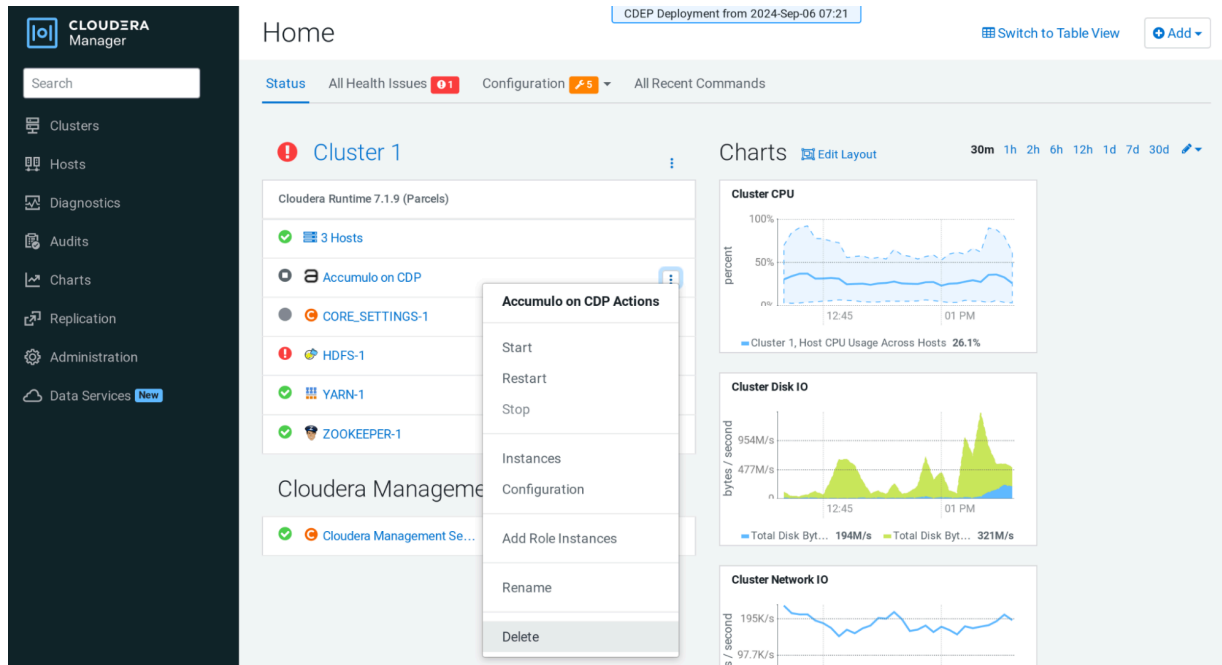
Note: You can use the Cloudera Manager REST API to export and import all of its configuration data. For more information, see [Backing Up and Restoring the Cloudera Manager Configuration](#).

7. Click Instances. Identify and note down the role assignments. This information is required when you add the Accumulo service to your cluster.

8. Click Cloudera Manager Clusters .

9.

Click  Delete .



Removing and updating Accumulo parcels

You must download, distribute, and activate the Accumulo 2.1.2 parcels.

Before you begin

You must install Accumulo CSD. For information, see [Install Operational Database powered by Apache Accumulo CSD file](#).

Procedure

1. Log in to Cloudera Manager Parcels .
2. Deactivate the Accumulo version 1.10.3
3. You must remove Accumulo from all the hosts and delete the Accumulo parcel.
4. Download the [Accumulo version 2.1.2 parcel](#).
5. Copy the ACCUMULO_ON_CDP-2.1.2_7.1.9.0.0.p0.11738502-el7.parcel file to the Local Parcel Repository. The Local Parcel Repository is by default located in /opt/cloudera/parcel-repo directory on the host where Cloudera Manager is running.
6. Click Distribute Activate .

Adding Accumulo service (unsecure)

You can add the Accumulo service (unsecure) on your CDP Private Cloud Base cluster.

Procedure

1.

Log in to Cloudera Manager Clusters  Add Service .

2. Select the Accumulo on CDP checkbox.




3. Assign the roles identified in the [Remove Accumulo 1.10.3](#) section.



Note: There are no Tracer roles in Accumulo 2.1.2.

4. Configure the Instance Name and Secret identified in the [Remove Accumulo 1.10.3](#) section.

5. Clear the Initialize Accumulo checkbox.

<p>Initialize Accumulo</p> <p>accumulo_initialize</p> <p> accumulo_initialize</p>	<p><input type="checkbox"/> Accumulo on CDP (Service-Wide) Undo i</p>
<p>Enable Kerberos authentication</p> <p>accumulo_secure_setup</p> <p> accumulo_secure_setup</p>	<p><input type="checkbox"/> Accumulo on CDP (Service-Wide) Undo i</p>
<p>Instance Volumes</p> <p>instance.volumes</p> <p> accumulo_instance_volumes</p>	<p>Accumulo on CDP (Service-Wide) i</p> <div style="border: 1px solid #ccc; height: 20px; width: 100%;"></div>

1 - 8 of 8

6. Specify Instance volumes.

This parameter specifies the filesystem URIs where Accumulo stores its data. It must be set when you want to explicitly control where Accumulo stores its data. If you do not explicitly set the instance.volumes option, Accumulo automatically uses the default filesystem URI from the Hadoop configuration.

7. Click Continue.

8. After the upgrade, reapply non-default configurations identified in the [Remove Accumulo 1.10](#) section. The Accumulo Service is now setup and running.

Adding Accumulo service (secure)

You can add the Accumulo service (secure) on your CDP Private Cloud Base cluster.

Before you begin

You must use Accumulo (secure) only if the previous Accumulo version was also secure.

Procedure

1.

Log in to Cloudera Manager Clusters  Add Service .

2. Select the Accumulo on CDP checkbox.

3. Assign the roles identified in the [Remove Accumulo 1.10.3](#) section.



Note: There are no Tracer roles in Accumulo 2.1.2.

4. Configure the Instance Name and Secret identified in the [Remove Accumulo 1.10.3](#) section.

5. Clear the Initialize Accumulo checkbox.

<p>Initialize Accumulo accumulo_initialize ⚙️ accumulo_initialize</p>	<p><input type="checkbox"/> Accumulo on CDP (Service-Wide) ↺ Undo</p>	<p>ⓘ</p>
<p>Enable Kerberos authentication accumulo_secure_setup ⚙️ accumulo_secure_setup</p>	<p><input type="checkbox"/> Accumulo on CDP (Service-Wide) ↺ Undo</p>	<p>ⓘ</p>
<p>Instance Volumes instance.volumes ⚙️ accumulo_instance_volumes</p>	<p>Accumulo on CDP (Service-Wide)</p> <div style="border: 1px solid #ccc; height: 20px; width: 100%;"></div>	<p>ⓘ</p>

1 - 8 of 8

6. Select Enable Kerberos Authentication.

7. Specify Instance volumes.

This parameter specifies the filesystem URIs where Accumulo stores its data. It must be set when you want to explicitly control where Accumulo stores its data. If you do not explicitly set the instance.volumes option, Accumulo automatically uses the default filesystem URI from the Hadoop configuration.

8. Click Continue.

9. After the upgrade, reapply non-default configurations identified in the [Remove Accumulo 1.10](#) section.

10. Create a client configuration file for the admin user.



Important: This step must be performed as a separate maintenance task. Zookeeper requires the same root user to be maintained across upgrades.

For more information, see the [Administrative User section of the user manual](#).

The Accumulo Service is now setup and running.

Upgrading Accumulo from 1.1.0 to 2.1.2

You can upgrade from Accumulo version 1.1.0 to Accumulo 2.1.2. Additionally, you must remove and reinstall the Accumulo service as a part of upgrading.

Upgrading CDP Private Cloud Base

You must upgrade your cluster from CDP Private Cloud Base 7.1.7.SP2 to CDP Private Cloud Base 7.1.9 SP1. For information on upgrading the CDP Private Cloud Base cluster, see [Upgrading CDP Private Cloud Base to a higher version](#).

Removing Accumulo 1.1.0

You must remove Accumulo 1.1.0 installed on your CDP cluster before upgrading the Accumulo service.

Before you begin

- You must stop all the ongoing administrative operations before you remove the Accumulo service.
- Make a note of all the non-default configurations before starting the upgrade process.

Procedure

1. Run the `root@accumulo> fate print` command to check for any Fault Tolerant Executor (FATE) transactions.
2. Run the `root@accumulo> fate delete <transaction id>` command to remove such transactions. Ensure tablet servers and the master are healthy.
3. Log in to Cloudera Manager Clusters Accumulo
4. Click Actions Stop .
5. Click Configuration.
6. Check all the parameters under the Configuration tab and identify and note down any non-default settings. This information is required when you add the Accumulo service to your cluster.



Note: 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](#).

7. Click Instances. Identify and note down the role assignments. This information is required when you add the Accumulo service to your cluster.
8. Click Cloudera Manager Clusters .
- 9.

Click  Delete .

Removing and updating Accumulo parcels

You must download, distribute, and activate the Accumulo 2.1.2 parcels.

Before you begin

You must install Accumulo CSD. For information, see [Install Operational Database powered by Apache Accumulo CSD file](#).

Procedure


1. Log in to Cloudera Manager Parcels .

2. Deactivate the Accumulo version 1.1.0
3. You must remove Accumulo from all the hosts and delete the Accumulo parcel.
4. Download the [Accumulo version 2.1.2 parcel](#).
5. Copy the ACCUMULO_ON_CDP-2.1.2_7.1.9.0.0.p0.11738502-el7.parcel file to the Local Parcel Repository. The Local Parcel Repository is by default located in /opt/cloudera/parcel-repo directory on the host where Cloudera Manager is running.
6. Click Distribute Activate .

Adding Accumulo service (unsecure)

You can add the Accumulo service (unsecure) on your CDP Private Cloud Base cluster.




Procedure

1. Log in to Cloudera Manager Clusters  Add Service .
2. Select the Accumulo on CDP checkbox.
3. Assign the roles identified in the [Remove Accumulo 1.1.0](#) section.



Note: There are no Tracer roles in Accumulo 2.1.2.

4. Configure the Instance Name and Secret identified in the [Remove Accumulo 1.1.0](#) section.
5. Clear the Initialize Accumulo checkbox.

<p>Initialize Accumulo accumulo_initialize  accumulo_initialize</p>	<p><input type="checkbox"/> Accumulo on CDP (Service-Wide) Undo ⓘ</p>
<p>Enable Kerberos authentication accumulo_secure_setup  accumulo_secure_setup</p>	<p><input type="checkbox"/> Accumulo on CDP (Service-Wide) Undo ⓘ</p>
<p>Instance Volumes instance.volumes  accumulo_instance_volumes</p>	<p>Accumulo on CDP (Service-Wide) ⓘ</p> <div style="border: 1px solid #ccc; height: 20px; width: 100%;"></div>

1 - 8 of 8

6. Specify Instance volumes.
This parameter specifies the filesystem URIs where Accumulo stores its data. It must be set when you want to explicitly control where Accumulo stores its data. If you do not explicitly set the instance.volumes option, Accumulo automatically uses the default filesystem URI from the Hadoop configuration.
7. Click Continue.
8. After the upgrade, reapply non-default configurations identified in the [Remove Accumulo 1.10](#) section. The Accumulo Service is now setup and running.

Adding Accumulo service (secure)

You can add the Accumulo service (secure) on your CDP Private Cloud Base cluster.

Before you begin

You must use Accumulo (secure) only if the previous Accumulo version was also secure.

Procedure

1.

Log in to Cloudera Manager Clusters  Add Service .


2. Select the Accumulo on CDP checkbox.

3. Assign the roles identified in the [Remove Accumulo 1.1.0](#) section.

Note: There are no Tracer roles in Accumulo 2.1.2.

4. Configure the Instance Name and Secret identified in the [Remove Accumulo 1.10](#) section.

5. Clear the Initialize Accumulo checkbox.

Initialize Accumulo <small>accumulo_initialize</small>  accumulo_initialize	<input checked="" type="checkbox"/> Accumulo on CDP (Service-Wide) Undo i
Enable Kerberos authentication <small>accumulo_secure_setup</small>  accumulo_secure_setup	<input type="checkbox"/> Accumulo on CDP (Service-Wide) Undo i
Instance Volumes <small>instance.volumes</small>  accumulo_instance_volumes	Accumulo on CDP (Service-Wide) i <input type="text"/>

1 - 8 of 8

6. Select Enable Kerberos Authentication.

7. Specify Instance volumes.

This parameter specifies the filesystem URIs where Accumulo stores its data. It must be set when you want to explicitly control where Accumulo stores its data. If you do not explicitly set the instance.volumes option, Accumulo automatically uses the default filesystem URI from the Hadoop configuration.

8. Click Continue.

9. After the upgrade, reapply non-default configurations identified in the [Remove Accumulo 1.10](#) section.

10. Create a client configuration file for the admin user.



Important: This step must be performed as a separate maintenance task. Zookeeper requires the same root user to be maintained across upgrades.

For more information, see the [Administrative User section of the user manual](#).

The Accumulo Service is now setup and running.