

Replication Manager for CDP Private Cloud Base

Date published: 2020-10-09

Date modified: 2020-11-25



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

Replication Manager Overview.....	5
Support matrix for Replication Manager on CDP Private Cloud Base.....	6
Port Requirements for Replication Manager.....	7
Data replication.....	8
Cloudera license requirements for Replication Manager.....	8
Replicating directories with thousands of files and subdirectories.....	8
Replication Manager log retention.....	8
Replicating from unsecure to secure clusters.....	9
Designating a replication source.....	9
Configuring a peer.....	10
Modifying peers.....	10
Configuring peers with SAML authentication.....	11
HDFS Replication.....	11
Source data.....	11
Network latency and replication.....	11
Performance and scalability limitations.....	11
HDFS replication from Sentry-enabled clusters.....	12
Guidelines for using snapshot diff-based replication.....	13
Configuring replication of HDFS data.....	14
Limiting replication hosts.....	18
Viewing replication policies.....	18
Viewing replication history.....	20
Monitoring the performance of HDFS replication policies.....	22
Hive/Impala replication.....	24
Host selection for Hive/Impala replication.....	25
Hive tables and DDL commands.....	25
Replication of parameters.....	25
Hive replication in dynamic environments.....	26
Creating a Hive/Impala replication policy.....	26
Sentry to Ranger replication for Hive replication policies.....	32
Replication of Impala and Hive User Defined Functions (UDFs).....	33
Monitoring the performance of Hive/Impala replication policies.....	33
Enabling, disabling, or deleting a replication policy.....	35

Replicating data to Impala clusters.....	36
Enabling replication between clusters with Kerberos Authentication.....	36
Ports.....	37
Considerations for realm names.....	37
HDFS, Hive, and Impala replication.....	37
Kerberos connectivity test.....	38
Copying data between a secure and an insecure cluster using DistCp and WebHDFS.....	39
Kerberos setup guidelines for Distcp between secure clusters.....	40
Replication of encrypted data.....	41
Encrypting data in transit between clusters.....	41
Security considerations.....	42
Snapshots.....	42
Cloudera Manager snapshot policies.....	42
Managing snapshot policies.....	43
Snapshots history.....	44
Orphaned snapshots.....	44
Managing HDFS snapshots.....	45
Browsing HDFS directories.....	45
Enabling and disabling HDFS snapshots.....	45
Taking and deleting HDFS snapshots.....	46
Restoring Snapshots.....	46
Using snapshots with replication.....	47
Hive/Impala replication using snapshots.....	47

Replication Manager Overview

Cloudera Manager provides an integrated, easy-to-use management solution to enable data protection on the Hadoop platform. Replication Manager enables you to replicate data across data centers for disaster recovery scenarios.

Replications can include data stored in HDFS, data stored in Hive tables, Hive metastore data, and Impala metadata (catalog server metadata) associated with Impala tables registered in the Hive metastore. When critical data is stored on HDFS, Cloudera Manager helps to ensure that the data is available at all times, even in case of complete shutdown of a data center. You can also use the HBase shell to replicate HBase data. (Cloudera Manager does not manage HBase replications.)

Replication Manager requires a valid license. To understand more about Cloudera license requirements, see [Managing Licenses](#).

Cloudera Manager provides the following key functionalities in the Cloudera Manager Admin Console:

- **Select** - Choose datasets that are critical for your business operations.
- **Schedule** - Create an appropriate schedule for data replication and snapshots. Trigger replication and snapshots as required for your business needs.
- **Monitor** - Track progress of your snapshots and replication jobs through a central console and easily identify issues or files that failed to be transferred.
- **Alert** - Issue alerts when a snapshot or replication job fails or is aborted so that the problem can be diagnosed quickly.

You can also use Cloudera Manager to schedule, save, and restore snapshots of HDFS directories and HBase tables.

Replication Manager functions consistently across HDFS and Hive:

- You can set it up on files or directories in HDFS and on external tables in Hive—without manual translation of Hive datasets to HDFS datasets, or vice versa. Hive Metastore information is also replicated.
- Applications that depend on external table definitions stored in Hive, operate on both replica and source as table definitions are updated.
- The hdfs user should have access to all Hive datasets, including all operations. Otherwise, Hive import fails during the replication process. To provide access, follow these steps:

1. Log in to Ranger Admin UI.
2. Navigate to the **Service Manager Hadoop_SQL Policies Access** section, and provide hdfs user permission to the all-database, table, column policy name.

Policy ID	Policy Name	Policy Labels	Status	Audit Logging	Roles	Groups	Users	Action
7	all - global	--	Enabled	Enabled	cdep-global_admin	--	rangerlookup, hive, beacon, dpprofiler + More...	View, Edit, Delete
8	all - database, table, column	--	Enabled	Enabled	cdep-global_admin	--	rangerlookup, hive, beacon, dpprofiler, hue, admin, impala, hdfs, OWNER - Less...	View, Edit, Delete
9	all - database, table	--	Enabled	Enabled	--	--	hive, beacon, dpprofiler, hue + More...	View, Edit, Delete
10	all - database	--	Enabled	Enabled	--	public	hive, beacon, dpprofiler, hue + More...	View, Edit, Delete
11	all - hiveservice	--	Enabled	Enabled	cdep-global_admin	--	rangerlookup, hive, beacon, dpprofiler + More...	View, Edit, Delete

You can also perform a “dry run” to verify configuration and understand the cost of the overall operation before actually copying the entire dataset.

Support matrix for Replication Manager on CDP Private Cloud Base

The support matrix contains compatibility information across features in Replication Manager. Replication Manager supports HDFS and Hive data replication. The matrix also lists the supported versions for CDH, CDP Private Cloud Base, and Cloudera Manager versions.

Feature	Lowest supported Cloudera Manager Version	Lowest supported CDH Version	Supported Services
Replication	Cloudera Manager 5.14+	CDH 5.13+	HDFS, Hive, Impala
Replication to and from Amazon S3*	Cloudera Manager 5.14+	CDH 5.13+	HDFS, Hive, Impala
Snapshots	Cloudera Manager 5.15+	CDH 5.15+	HDFS, Hive, Impala
Replication to and from Microsoft ADLS Gen1	Cloudera Manager 5.15, 5.16, 6.1+	CDH 5.13+	HDFS, Hive, Impala
Replication to Microsoft ADLS Gen2 (ABFS)	Cloudera Manager 6.1+	CDH 5.13+	HDFS, Hive, Impala

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



Note: Replication Manager service is supported from CDP Private Cloud Base version 7.0.3 and Cloudera Manager version 7.0.3 onwards.

Starting in Cloudera Manager 6.1.0, Replication Manager ignores Hive tables backed by Kudu during replication. The change does not affect functionality since Replication Manager does not support tables backed by Kudu. This change was made to guard against data loss due to how the Hive Metastore, Impala, and Kudu interact.

Supported replication scenarios

Kerberos

Replication Manager supports the following replication scenarios when Kerberos authentication is used on a cluster:

- Secure source to a secure destination.
- Insecure source to an insecure destination.
- Insecure source to a secure destination. The following requirements must be met for this scenario:
 - When a destination cluster has multiple source clusters, all the source clusters must either be secure or insecure. Replication Manager does not support a mix of secure and insecure source clusters.
 - The destination cluster must run Cloudera Manager 7.x or higher.
 - The source cluster must run a compatible Cloudera Manager version.
 - This replication scenario requires additional configuration. For more information, see [Replicating from unsecure to secure clusters](#) on page 9.

Cloud storage

Replication Manager supports replicating to or from Amazon S3, Microsoft Azure ADLS Gen1, and Microsoft Azure ADLS Gen2 (ABFS).

Transport Layer Security (TLS)

You can use TLS with Replication Manager. Additionally, Replication Manager supports replication scenarios where TLS is enabled for non-Hadoop services (Hive/Impala) and TLS is disabled Hadoop services (such as HDFS, YARN, and MapReduce).

Sentry-related replication

To perform Apache Sentry to Apache Ranger replication using HDFS and Hive replication policies, you must have installed Cloudera Manager version 6.3.1 and higher on the source cluster and Cloudera Manager version 7.1.1 and higher on the target cluster.

When the source cluster is Sentry-enabled and you want to run HDFS replication policies, use the `hdfs` user to run the replication policy. The replication policy copies the permissions of replicated files and tables to the target cluster. To use any other user account, make sure that you configure the user account to bypass Sentry ACLs during replication.

When you create a Hive replication policy, choose the appropriate options to ensure that the Sentry permissions are migrated to Ranger permissions. The Replication Manager uses the `authzmigrator` tool to move data from Sentry to Ranger during Hive replication.

Unsupported Replication Scenarios**Versions**

Replicating to and from HDP to Cloudera Manager 7.x is not supported.

Hive replication

Hive replication policies do not support managed to managed table replication. When you replicate from a CDH cluster to a CDP Private Cloud Base cluster, Replication Manager converts managed tables to external tables.

Kerberos

When Kerberos authentication is used on a cluster, replication from a secure source to an insecure destination is not supported.

Ranger

Ranger to Ranger replication is not supported.

Apache Knox

If Cloudera Manager is configured with Knox, Replication Manager does not work.

Port Requirements for Replication Manager

Make sure that the following ports are open and accessible on the source hosts to allow communication between the source and destination Cloudera Manager servers and the HDFS, Hive, MapReduce, and YARN hosts.

Port details for Replication Manager

Service	Default Port
Cloudera Manager Admin Console HTTP	7180
Cloudera Manager Admin Console HTTPS (with TLS enabled)	7183
Cloudera Manager Agent	9000
HDFS NameNode	8020
Ranger KMS	9292 9494 if TLS enabled
HDFS DataNode	9866
NameNode WebHDFS	9870
YARN Resource Manager	8032
DataNode Secure	1004
NameNode Secure WebHDFS	9871

Service	Default Port
Hive Metastore	9083
Impala Catalog Server	26000

Data replication

Before you use Replication Manager, you must understand some of the requirements about data replication.

Cloudera license requirements for Replication Manager

You must have the license to perform your tasks in Replication Manager

To understand more about Cloudera license requirements, see [Managing Licenses](#).

Replicating directories with thousands of files and subdirectories

Before you replicate the data in directories that has thousands of files and subdirectories, increase the heap size in the `hadoop-env.sh` file.

Procedure

1. To increase the heap size, go to the HDFS service page on the destination Cloudera Manager instance.
2. Click the Configuration tab.
3. Expand SCOPE and select HDFS service name (Service-Wide) option.
4. Expand CATEGORY and select Advanced.
5. Locate the HDFS Replication Environment Advanced Configuration Snippet (Safety Valve) for `hadoop-env.sh` property.
6. To increase the heap size, add the key-value pair `HADOOP_CLIENT_OPTS=-Xmx<memory_value>`. For example, if you enter `HADOOP_CLIENT_OPTS=-Xmx1g`, the heap size is set to 1 GB. This value should be adjusted depending on the number of files and directories being replicated.
7. Enter a Reason for change, and then click Save Changes to commit the changes.

Replication Manager log retention

By default, Cloudera Manager retains Replication Manager logs for 90 days. You can change the number of days Cloudera Manager retains logs or disable log retention completely.

1. Navigate to the Cloudera Manager *HDFS Service* Configuration section.
2. In the Cloudera Manager Admin Console, search for the Replication Manager Log Retention property.
3. Enter the number of days you want to retain the logs.
4. To disable log retention, enter -1.



Important: Automatic log expiration purges custom set replication log and metadata files too. These paths are set by Log Path and Directory for Metadata arguments that are present on the UI as per the schedule fields. It is the user's responsibility to set valid paths (For example, specify the legal HDFS paths that are writable by current user) and maintain this information for each replication policy.

Replicating from unsecure to secure clusters

You can use Replication Manager to replicate data from an unsecure cluster, one that does not use Kerberos authentication, to a secure cluster, a cluster that uses Kerberos. Note that the reverse is not true.

About this task

Replication Manager does not support replicating from a secure cluster to an unsecure cluster. To perform the replication, the destination cluster must be managed by Cloudera Manager 6.1.0 or higher. The source cluster must run Cloudera Manager 5.14.0 or higher in order to be able to replicate to Cloudera Manager 6.



Note: In replication scenarios where a destination cluster has multiple source clusters, all the source clusters must either be secure or unsecure. Replication Manager does not support replication from a mixture of secure and unsecure source clusters.

To enable replication from an unsecure cluster to a secure cluster, you need a user that exists on all the hosts on both the source cluster and destination cluster. Specify this user in the Run As Username field when you create a replication policy.

Procedure

1. On a host in the source or destination cluster, add a user with the following command:

```
sudo -u hdfs hdfs dfs -mkdir -p /user/<username>
```

For example, the following command creates a user named milton:

```
sudo -u hdfs hdfs dfs -mkdir -p /user/milton
```
2. Set the permissions for the user directory with the following command:

```
sudo -u hdfs hdfs dfs -chown <username> /user/username
```

For example, the following command makes milton the owner of the milton directory:

```
sudo -u hdfs hdfs dfs -chown milton /user/milton
```
3. Create the supergroup group for the user you created in step 1 with the following command:

```
groupadd supergroup
```
4. Add the user you created in step 1 to the group you created:

```
usermod -G supergroup <username>
```

For example, add milton to the group named supergroup:

```
usermod -G supergroup milton
```
5. Repeat this process for all hosts in the source and destination clusters so that the user and group exists on all of them.

What to do next

After you complete this process, specify the user you created in the Run As Username field when you create a replication policy.

Designating a replication source

You must assign the source cluster to replicate the data.

Minimum Required Role: [Cluster Administrator](#) (also provided by Full Administrator).

The Cloudera Manager Server that you are logged into is the destination for replications set up using that Cloudera Manager instance. From the Admin Console of this destination Cloudera Manager instance, you can designate a peer Cloudera Manager Server as a source of HDFS and Apache Hive data for replication.

Configuring a peer

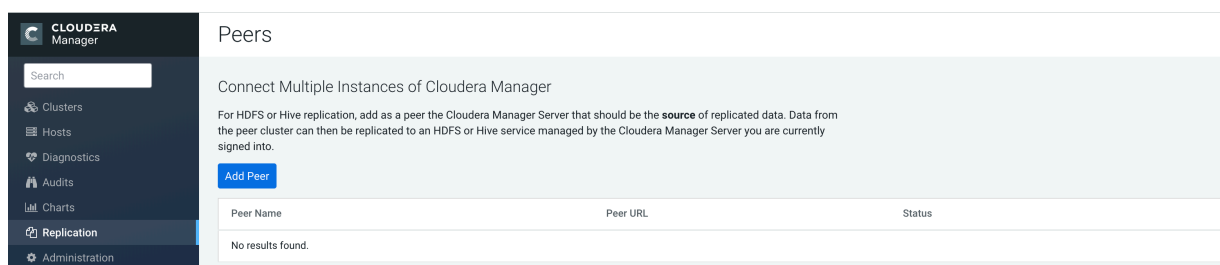
Before you replicate data from source cluster to destination cluster, you must connect the Cloudera Manager with the peer and then test the connectivity.

About this task

If your cluster uses SAML Authentication, see [Configuring peers with SAML authentication](#) on page 11 before configuring a peer.

Procedure

1. From Cloudera Manager, select **Replication Peers** . If there are no existing peers, a **Add Peer** button appears along with a short message. If peers already exist, they appear in the **Peers** list.



2. Click **Add Peer**.
3. In the **Add Peer** dialog box, provide a name, the peer URL (including the port) of the Cloudera Manager Server source for the data to be replicated, and the login credentials for that server.



Important: The role assigned to the login on the source server must be either a *User Administrator* or a *Full Administrator*.

Cloudera recommends that TLS/SSL be used. A warning is shown if the URL scheme is http instead of https. After configuring both peers to use TLS/SSL, add the remote source Cloudera Manager TLS/SSL certificate to the local Cloudera Manager truststore, and vice versa.

4. Click the **Add** button in the dialog box to create the peer relationship.

Results

The peer is added to the **Peers** list. Cloudera Manager automatically tests the connection between the Cloudera Manager Server and the peer. You can also click **Test Connectivity** to test the connection. **Test Connectivity** also tests the Kerberos configuration for the clusters.

Modifying peers

You can modify or delete peers.

Procedure

1. To edit a peer, select a peer and click **Actions Edit** .
2. Make your changes.
3. Click **Update Peer** to save your changes.
4. To delete a peer, select a peer and click **Actions Delete** ..

Configuring peers with SAML authentication

If your cluster uses SAML Authentication, you can create a Cloudera Manager user account that has the User Administrator or Full Administrator role before you create a peer.



Note: You can also use an existing user that has one of these roles. You can use this user to create the peer relationship, and then optionally, you can delete the user account.

HDFS Replication

Replication related to HDFS data is discussed in this section.

Minimum Required Role: [Replication Administrator](#) (also provided by Full Administrator)

HDFS replication enables you to copy (replicate) your HDFS data from one HDFS service to another, synchronizing the data set on the destination service with the data set on the source service, based on a specified replication policy. The destination service must be managed by the Cloudera Manager Server where the replication is being set up, and the source service can be managed by that same server or by a peer Cloudera Manager Server. You can also replicate HDFS data within a cluster by specifying different source and destination directories.

Remote Replication Manager automatically copies HDFS metadata to the destination cluster as it copies files. HDFS metadata need only be backed up locally.

Source data

Before you start a replication job on a directory, all the files in the directory must be closed. Replication fails if source files are open. If you cannot ensure that all source files are closed, you can configure the replication to continue despite errors by clearing the Abort on Error option for HDFS replication in the replication wizard. When a replication job runs, you must make sure that the source directory is not modified. A file added during replication does not get replicated. If you delete a file during replication, the replication fails.

After the replication completes, you can view the log for the replication to identify opened files. Ensure these files are closed before the next replication occurs.

Network latency and replication

High latency among clusters can cause replication jobs to run more slowly, but does not cause them to fail.

For best performance, latency between the source cluster NameNode and the destination cluster NameNode should be less than 80 milliseconds. You can test latency using the Linux ping command. Cloudera has successfully tested replications with latency of up to 360 milliseconds. As latency increases, replication performance degrades.

Performance and scalability limitations

HDFS replication has some limitations.

- The maximum number of files for a single replication job is 100 million.
- The maximum number of files for a replication policy that runs more frequently than once in 8 hours is 10 million.
- The throughput of the replication job depends on the absolute read and write throughput of the source and destination clusters.
- Regular rebalancing of your HDFS clusters is required for efficient operation of replications.



Note: Cloudera Manager provides downloadable data that you can use to diagnose HDFS replication performance.

HDFS replication from Sentry-enabled clusters

When you run a HDFS replication policy on a Sentry-enabled source cluster, the replication policy can copy files and tables along with their permissions.

Before you begin

Cloudera Manager version 6.3.1 and above is required to run HDFS replication policies on a Sentry-enabled source cluster.

When you want to run HDFS replication policies on a source cluster that is Sentry-enabled, you must use the `hdfs` user. If you want to use a different user account, you must configure the user account to bypass the Sentry ACLs during the replication process.

When Sentry is not available or when Sentry does not manage the authorization for a resource such file or directory in the source cluster, HDFS uses its internal ACLs to manage resource authorization.

When Sentry is enabled for the source cluster and you use the `hdfs` user name to run the HDFS replication policy, HDFS copies the ACLs configured in Sentry for the replicated files and tables to the target cluster.

When Sentry is enabled and you use a different user name to run the HDFS replication policy, both Sentry ACLs and HDFS internal ACLs are copied which results in incorrect HDFS metadata in the target cluster. If the Sentry ACLs are not compatible with HDFS ACLs, the replication job fails.

To avoid compatibility issues between HDFS and Sentry ACLs for a non-`hdfs` user, you must complete the following steps:

Procedure

1. Create a user account that is only used for Replication Manager jobs since Sentry ACLs will be bypassed for this user.

For example, create a user named `bdr-only-user`.

2. To bypass the Sentry ACLs during replication, perform the following steps on the source cluster:

- a) In the Cloudera Manager Admin Console, select *Clusters* *HDFS service*.
- b) Select Configuration and search for NameNode Advanced Configuration Snippet (Safety Valve) for `hdfs-site.xml` property.
- c) Add the following property:

Name - `dfs.namenode.inode.attributes.provider.bypass.users`

Value - Enter `[***USERNAME, USERNAME@REALMNAME***]`, where `[***USERNAME***]` is the user you created in step 1 and the `[***REALMNAME***]` is the Kerberos realm name.

For example, if the username is `bdr-only-user` on the realm `elephant`, enter **`bdr-only-user, bdr-only-user@ElephantRealm`**

- d) Restart the NameNode.

3. Repeat step 2 on the destination cluster.
4. When you create a HDFS replication policy, specify the user you created in step 1 in the Run As Username and Run on Peer as Username (if available) fields.



Note: The Run As Username field is used to launch MapReduce job for copying data. Run on Peer as Username field is used to run copy listing on source, if different than Run as Username.

What to do next



Important:

When the source cluster is Sentry-enabled, make sure to use the `hdfs` user name or the configured username as Run on Peer as Username and Run as Username. If you use different usernames, the Replication Manager reads the ACLs from the source as `hdfs`, replicates the Sentry ACLs to the target cluster, and applies them to the files in HDFS. This results in additional usage of NameNode heap in the target cluster.

Guidelines for using snapshot diff-based replication

By default, Replication Manager uses snapshot differences ("diff") to improve performance by comparing HDFS snapshots and only replicating the files that are changed in the source directory. While Hive metadata requires a full replication, the data stored in Hive tables can take advantage of snapshot diff-based replication.

After every replication, the Replication Manager retains a snapshot on the source cluster. Using the snapshot copy on the source cluster, Replication Manager performs incremental backup for the next replication cycle.

Replication Manager retains snapshots on the source cluster and uses snapshot diff-based replication only if:

- Source and target clusters are managed by Cloudera Manager 5.15 and higher.
- The source cluster must be managed by Cloudera Manager 5.15.0 or higher when the destination is Amazon S3 or Microsoft ADLS.



Important: Snapshot-diff-based replication from S3/ABFS to HDFS is not supported because S3/AABFS do not support snapshots.

- Source and target CDH versions are 5.13.3 or higher, 5.14.2 or higher, and 5.15 or higher.

You must follow the below guidelines to use snapshot diff-based replication efficiently in replication policies:

- The source and target clusters must be managed by Cloudera Manager 5.15.0 or higher.
- The source and target clusters must run CDH version 5.15.0 or higher, 5.14.2 or higher, or 5.13.3 or higher.
- The HDFS snapshots must be immutable.



Tip: In Cloudera Manager, go to the *Clusters HDFS service Configuration* section, and search for *Enable Immutable Snapshots*.

- The snapshot root directory must be set as low in the hierarchy as possible.
- To run the job, the user must be a super user or the owner of the snapshottable root. This is because the run-as-user (specified in the replication policy) must have the required permissions to list the snapshots.
- The paths from both source and destination clusters in the replication policy must be present under a snapshottable root, or must be snapshottable.



Tip: An HDFS directory is referred to as snapshottable if an administrator - having superuser privilege or having owner access to the directory - has enabled snapshots for the directory in Cloudera Manager.

- All the HDFS paths for the tables in a database must be snapshottable or under a snapshottable root for a Hive replication policy to replicate the database successfully.

For example, if the database being replicated has external tables, all the external table HDFS data locations should be snapshottable. This is because if the external table locations are not snapshottable, Replication Manager does not generate a diff report. The Replication Manager needs a diff report to use the snapshot diff feature.



Important: Do not use snapshot diff for globbed paths because it is not optimized for globbed paths.

FAQs

What do I do when snapshot diff-based replication fails because an encrypted subdirectory exists in the source data?

To resolve this issue, create an exclusion regex in the replication policy to exclude the subdirectory during replication. Create another replication policy to replicate the encrypted subdirectory.

During what circumstances does the Replication Manager initiate a complete data replication?

Replication Manager initiates a complete replication for the following scenarios:

- When you do not choose Abort on Snapshot Diff Failures (when you create a replication policy in Replication Manager) and errors appear during the replication process.

In this case, the Replication Manager continues to replicate and performs a complete replication after it encounters an error.

- When one or more of the following parameters that you set in the replication policy changes:
 - Delete Policy
 - Preserve Policy
 - Target Path
 - Exclusion Path.
- When a change in the target directories is detected.

Replication Manager ensures that the next HDFS snapshot replication is a complete replication.

Configuring replication of HDFS data

You must set up your clusters before you plan HDFS data replication.

Procedure

1. Verify that your cluster conforms to one of the supported replication scenarios.
2. If you are using different Kerberos principals for the source and destination clusters, add the destination principal as a proxy user on the source cluster. For example, if you are using the hdfssrc principal on the source cluster and the hdfsdest principal on the destination cluster, add the following properties to the HDFS service Cluster-wide Advanced Configuration Snippet (Safety Valve) for core-site.xml property on the source cluster:

```
<property>
  <name>hadoop.proxyuser.hdfsdest.groups</name>
  <value>*</value>
</property>
<property>
  <name>hadoop.proxyuser.hdfsdest.hosts</name>
  <value>*</value>
</property>
```

Deploy the client configuration and restart all services on the source cluster, if the source cluster is managed by a different Cloudera Manager server than the destination cluster.

3. From Cloudera Manager > Replication > Replication Policies page, click Create Replication Policy.

Replication Policies

Search

Filters

STATUS

- Failed 0
- Succeeded 0
- Running 0
- Disabled 0
- Dry-run 0

TYPE

- HDFS 0
- HDFS-S3 0
- Hive 0
- Hive-S3 0
- HDFS-ADLS 0
- Hive-ADLS 0

Actions for Selected

Create Replication Policy

ID	Name	Type	Source	Destination	Throughput	Progress	Completed	Next Run
No replication policies.								

Last Refreshed 6:04 AM

4. Select HDFS Replication Policy.

Replication Policies

Search

Filters

STATUS

- Failed 0
- Succeeded 0
- Running 0
- Disabled 0
- Dry-run 0

TYPE

- HDFS 0
- HDFS-S3 0
- Hive 0
- Hive-S3 0
- HDFS-ADLS 0
- Hive-ADLS 0

Actions for Selected

Create Replication Policy

HDFS Replication Policy

Hive Replication Policy

ID	Name	Type	Source	Destination	Throughput	Progress	Completed	Next Run
No replication policies.								

Last Refreshed 6:36 AM

The **Create HDFS Replication Policy** dialog box appears.

5. In the **General** tab, you can configure the following options:

- Click the Name field and add a unique name for the replication policy.
- Click the Source field and select the source HDFS service. You can select HDFS services managed by a peer Cloudera Manager Server, local HDFS services (managed by the Cloudera Manager Server for the Admin Console you are logged into).
- Enter the Source Path to the directory (or file) you want to replicate.
- Click the Destination field and select the destination HDFS service from the HDFS services managed by the Cloudera Manager Server for the Admin Console you are logged into.
- Enter the Destination Path where the source files should be saved.
- Select a Schedule:
 - Immediate - Run the schedule Immediately.
 - Once - Run the schedule one time in the future. Set the date and time.
 - Recurring - Run the schedule periodically in the future. Set the date, time, and interval between runs.
- Enter the user to run the replication job in the Run As Username field. By default this is hdfs. If you want to run the job as a different user, enter the user name here. If you are using Kerberos, you must provide a user name here, and it must be one with an ID greater than 1000. (You can also configure the minimum user ID

number with the `min.user.id` property in the YARN or MapReduce service.) Verify that the user running the job has a home directory, `/user/username`, owned by `username:supergroup` in HDFS. This user must have permissions to read from the source directory and write to the destination directory.

Note the following:

- The User must not be present in the list of banned users specified with the Banned System Users property in the YARN configuration (Go to the YARN service, select Configuration tab and search for the property). For security purposes, the `hdfs` user is banned by default from running YARN containers.
- The requirement for a user ID that is greater than 1000 can be overridden by adding the user to the "white list" of users that is specified with the Allowed System Users property. (Go to the YARN service, select the Configuration tab and search for the property.)

6. Select the Resources tab to configure the following options:

- Scheduler Pool – (Optional) Enter the name of a resource pool in the field. The value you enter is used by the MapReduce Service you specified when Cloudera Manager executes the MapReduce job for the replication. The job specifies the value using one of these properties:
 - MapReduce – Fair scheduler: `mapred.fairscheduler.pool`
 - MapReduce – Capacity scheduler: `queue.name`
 - YARN – `mapreduce.job.queue.name`
- Maximum Map Slots - Limits for the number of map slots per mapper. The default value is 20.
- Maximum Bandwidth - Limits for the bandwidth per mapper. The default is 100 MB.
- Replication Strategy - Whether file replication tasks should be distributed among the mappers statically or dynamically. (The default is Dynamic.) Static replication distributes file replication tasks among the mappers up front to achieve a uniform distribution based on the file sizes. Dynamic replication distributes file replication tasks in small sets to the mappers, and as each mapper completes its tasks, it dynamically acquires and processes the next unallocated set of tasks.

7. Select the Advanced Options tab to configure the following options:

- Add Exclusion - Click the link to exclude one or more paths from the replication. The Regular Expression-Based Path Exclusion field displays, where you can enter a regular expression-based path. When you add an

exclusion, include the snapshotted relative path for the regex. For example, to exclude the /user/bdr directory, use the following regular expression, which includes the snapshots for the bdr directory:

```
.*/user/\.snapshot/./+/bdr.*
```

To exclude top-level directories from replication in a globbed source path, you can specify the relative path for the regex without including .snapshot in the path. For example, to exclude the bdr directory from replication, use the following regular expression:

```
.*/user+/bdr.*
```

You can add more than one regular expression to exclude.

- MapReduce Service - The MapReduce or YARN service to use.
- Log path - An alternate path for the logs.
- Description - A description of the replication policy.
- Error Handling You can select the following:
 - Skip Checksum Checks - Whether to skip checksum checks on the copied files. If checked, checksums are not validated. Checksums are checked by default.



Important: You must skip checksum checks to prevent replication failure due to non-matching checksums in the following cases:

- Replications from an encrypted zone on the source cluster to an encrypted zone on a destination cluster.
- Replications from an encryption zone on the source cluster to an unencrypted zone on the destination cluster.
- Replications from an unencrypted zone on the source cluster to an encrypted zone on the destination cluster.

Checksums are used for two purposes:

- To skip replication of files that have already been copied. If Skip Checksum Checks is selected, the replication job skips copying a file if the file lengths and modification times are identical between the source and destination clusters. Otherwise, the job copies the file from the source to the destination.
- To redundantly verify the integrity of data. However, checksums are not required to guarantee accurate transfers between clusters. HDFS data transfers are protected by checksums during transfer and storage hardware also uses checksums to ensure that data is accurately stored. These two mechanisms work together to validate the integrity of the copied data.
- Skip Listing Checksum Checks - Whether to skip checksum check when comparing two files to determine whether they are same or not. If skipped, the file size and last modified time are used to determine if files are the same or not. Skipping the check improves performance during the mapper phase. Note that if you select the Skip Checksum Checks option, this check is also skipped.
- Abort on Error - Whether to abort the job on an error. If selected, files copied up to that point remain on the destination, but no additional files are copied. Abort on Error is off by default.
- Abort on Snapshot Diff Failures - If a snapshot diff fails during replication, Replication Manager uses a complete copy to replicate data. If you select this option, the Replication Manager aborts the replication when it encounters an error instead.
- Preserve - Whether to preserve the block size, replication count, permissions (including ACLs), and extended attributes (XAttrs) as they exist on the source file system, or to use the settings as configured on the destination file system. By default source system settings are preserved. When Permission is checked, and both the source and destination clusters support ACLs, replication preserves ACLs. Otherwise, ACLs are not replicated. When Extended attributes is checked, and both the source and destination clusters support extended attributes,

replication preserves them. (This option only displays when both source and destination clusters support extended attributes.)



Note: To preserve permissions to HDFS, you must be running as a superuser on the destination cluster. Use the Run As Username option to ensure that is the case.

- **Delete Policy** - Whether files that were deleted on the source should also be deleted from the destination directory. This policy also determines the handling of files in the destination location that are unrelated to the source. Options include:
 - **Keep Deleted Files** - Retains the destination files even when they no longer exist at the source. (This is the default.).
 - **Delete to Trash** - If the HDFS trash is enabled, files are moved to the trash folder.
 - **Delete Permanently** - Uses the least amount of space; use with caution. This option does not delete the files and directories in the top level directory. This is in line with rsync/Hadoop DistCp behaviour.
- **Alerts** - Whether to generate alerts for various state changes in the replication workflow. You can alert on failure, on start, on success, or when the replication workflow is aborted.

8. Click Save Policy.

The replication task now appears as a row in the **Replication Policies** table. It can take up to 15 seconds for the task to appear.

If you selected Immediate in the Schedule field, the replication job begins running when you click Save Policy.

What to do next

To specify additional replication tasks, select **Create HDFS Replication**.



Note: If your replication job takes a long time to complete, and files change before the replication finishes, the replication may fail. Consider making the directories snapshottable, so that the replication job creates snapshots of the directories before copying the files and then copies files from these snapshottable directories when executing the replication.

Limiting replication hosts

If your cluster has clients installed on hosts with limited resources, HDFS replication may use these hosts to run commands for the replication, which can cause performance degradation. You can limit HDFS replication to run only on selected DataNodes by specifying a "whitelist" of DataNode hosts.

Procedure

1. Click **Clusters** *HDFS service* Configuration .
2. Type HDFS Replication in the search box.
3. Locate the HDFS Replication Environment Advanced Configuration Snippet (Safety Valve) property.
4. Add the HOST_WHITELIST property. Enter a comma-separated list of DataNode hostnames to use for HDFS replication. For example:

```
HOST_WHITELIST=host-1.mycompany.com,host-2.mycompany.com
```

5. Click **Save Changes** to commit the changes.

Viewing replication policies

The Replications Policies page displays a row of information about each replication policy. Each row also displays recent messages regarding the last time the replication job ran.

Figure 1: Replication Policies

Replication Policies

Search

Filters

STATUS

- Failed 1
- Succeeded 2
- Running 0
- Disabled 0
- Dry-run 0

TYPE

- HDFS 2
- HDFS-S3 0
- Hive 1
- Hive-S3 0
- HDFS-ADLS 0
- Hive-ADLS 0

SOURCE

- Cluster 1 1
- Cluster 1 @ test 2

TARGET

Actions for Selected

Create Replication Policy

ID	Name	Type	Source	Destination	Throughput	Progress	Completed	Next Run
5	test	HDFS	HDFS-1 Cluster 1	HDFS-1 Cluster 1			7:59 PM	None scheduled.
Message: HDFS replication command succeeded. From: /tmp To: /tmp/rece								
8	tsadf	HDFS	HDFS-1 Cluster 1 @ test	HDFS-1 Cluster 1			12:33 AM	None scheduled.
Message: HDFS replication command succeeded. From: /tmp To: /tmp/rec								
12	testadsf	Hive	HIVE-1 Cluster 1 @ test	HIVE-1 Cluster 1			1:29 AM	None scheduled.
Message: Hive Replication Import step failed. Objects: Custom Databases								

Only one job corresponding to a replication policy can occur at a time; if another job associated with that same replication policy starts before the previous one has finished, the second one is canceled.

You can limit the replication jobs that are displayed by selecting filters on the left. If you do not see an expected policy, adjust or clear the filters. Use the search box to search the list of policies for path, database, or table names.

The Replication Policies columns are described in the following table:

Table 1: Replication Policies Table

Column	Description
ID	An internally generated ID number that identifies the policy. Provides a convenient way to identify a policy. Click the ID column label to sort the replication policies table by ID.
Name	The unique name you specify when you create a policy. Click the Name column label to sort the replication policies table by name.
Type	The type of replication policy, HDFS or Hive.
Source	The source cluster for the replication.
Destination	The destination cluster for the replication.
Throughput	Average throughput per mapper/file of all the files written. Note that throughput does not include the following information: the combined throughput of all mappers and the time taken to perform a checksum on a file after the file is written.
Progress	The progress of the replication.
Completed	The time when the replication job completed. Click the Completed column label to sort the replication policies table by time.
Next Run	The date and time when the next replication is scheduled, based on the schedule parameters specified for the policy. Hover over the date to view additional details about the scheduled replication. Click the Next Run column label to sort the replication policies table by the next run date.

Column	Description
Actions	<p>The following items are available from the Action button:</p> <ul style="list-style-type: none"> • Show History. Opens the Replication History page for a replication. • Edit Configuration. Opens the Edit HDFS Replication Policy page. • Dry Run. Simulates a run of the replication task but does not actually copy any files or tables. After a Dry Run, you can select Show History, which opens the Replication History page where you can view any error messages and the number and size of files or tables that would be copied in an actual replication. • Run Now - Runs the replication task immediately. • Collect Diagnostic Data. Opens the Send Diagnostic Data screen, which allows you to collect replication-specific diagnostic data for the last 10 runs of the policy. <p>In the Send Diagnostic Data screen, select Send Diagnostic Data to Cloudera to automatically send the bundle to Cloudera Support. You can also enter a ticket number and comments when sending the bundle. After you click Collect and Send Diagnostic Data, the Replication Manager generates the bundle and opens the Replications Diagnostics Command screen. When the command finishes, click Download Result Data to download a zip file containing the bundle.</p> <ul style="list-style-type: none"> • Disable Enable. Disables or enables the replication policy. No further replications are scheduled for disabled replication policies. • Delete. Deletes the policy. Deleting a replication policy does not delete copied files or tables.

Viewing replication history

You can view the historical details about replication jobs on the Replication History page.

To view the history of a replication job:

1. From Cloudera Manager, select Replication > Replication Policies.

The list of available replication policies appear.

2. Locate the row for the policy, select the policy, and click Actions. Select Show History.

The Replication History page appears with the job information.

Figure 2: Replication History Screen (HDFS)

Replication Policies

Replication History											
Name	test	Type	HDFS	Source	HDFS-1 (Cluster 1)	Destination	HDFS-1 (Cluster 1)	Next Run	None scheduled.		
Start Time	Duration	Outcome	Files Expected		Files Copied		Files Failed		Files Deleted	Files Skipped	
▼ September 23, 2020 7:58 PM	1 min	Successful	80	(722.5 MiB)	17	(94.4 MiB)	0	(0 B)	0	63	(628.1 MiB)
Started At	September 23, 2020 7:58 PM										
Duration	a few seconds										
Command Details	View										
MapReduce Job	job_1600880827337_0009										
HDFS Replication Report	Download CSV										
Message	17 file(s) copied, 63 unchanged.										

Replication History Table

The Replication History page displays a table of previously run replication jobs with the following columns:

Column	Description
Start Time	<p>Shows the details about the job.</p> <p>You can expand the section to view the following job details:</p> <ul style="list-style-type: none"> Started At - Displays the time the replication job started. Duration - Displays the time duration for the job to complete. Command Details - Displays the command details in a new tab after you click View. <p>The Command Details page displays the details and messages about each step during command run. On this page, click Context to view the service status page relevant to the command, and click Download to download the summary as a JSON file.</p> <p>To view the command details, expand the Step section and then choose Show All Steps, Show Only Failed Steps, or Show Only Running Steps. In this section, you can perform the following tasks:</p> <ul style="list-style-type: none"> View the actual command string. View the start time and duration for the command run. View the host status page for the command by clicking the host link. View the full log file for the command by selecting the stdout or stderr tab. <p>See Viewing Running and Recent Commands.</p> <ul style="list-style-type: none"> MapReduce Job. Click the link to view the job details. HDS Replication Report. Click Download CSV to view the following options: <ul style="list-style-type: none"> Listing - Click to download the CSV file that contains the replication report. The file lists the list of files and directories copied during the replication job. Status - Click to download the CSV file that contains the complete status report. The file contains the full status report of the files where the status of the replication is one of the following: <ul style="list-style-type: none"> ERROR – An error occurred and the file was not copied. DELETED – A deleted file. SKIPPED – A file where the replication was skipped because it was up-to-date. Error Status Only - Click to download the CSV file that contains the status report of all copied files with errors. The file lists the status, path, and message for the copied files with errors. Deleted Status Only - Click to download the CSV file that contains the status report of all deleted files. The file lists the status, path, and message for the databases and tables that were deleted. Skipped Status Only - Click to download the CSV file that contains the status report of all skipped files. The file lists the status, path, and message for the databases and tables that were skipped. Performance - Click to download a CSV file which contains a summary report about the performance of the running replication job. The performance summary report includes the last performance sample for each mapper that is working on the replication job. Full Performance - Click to download the CSV file that contains the performance report of the job. The performance report shows the samples taken for all the mappers during the full execution of the replication job. (Dry Run only) View the number of Replicable Files. Displays the number of files that would be replicated during an actual replication. (Dry Run only) View the number of Replicable Bytes. Displays the number of bytes that would be replicated during an actual replication. View the number of Impala UDFs replicated. (Displays only for Hive/Impala replications where Replicate Impala Metadata is selected.) If a user was specified in the Run As Username field when creating the replication job, the selected user displays. View messages returned from the replication job.
Duration	Time taken for the replication job to complete.
Outcome	Indicates the status of the replication job as Successful or Failed.
Files Expected	Number of files expected to be copied and its file size based on the parameters of the replication policy.
Files Copied	Number of files copied and its file size for the replication job.
Files Failed	Number of files that failed to be copied and its file size for the replication job.
Files Deleted	Number of files that were deleted and its file size for the replication job
Files Skipped	Number of files skipped and its file size for the replication job. The replication process skips files that already exist in the destination and have not changed.

Monitoring the performance of HDFS replication policies

You can monitor the progress of an HDFS replication policy using performance data that you download as a CSV file from the Cloudera Manager Admin console.

About this task

This file contains information about the files being replicated, the average throughput, and other details that can help diagnose performance issues during HDFS replications. You can view this performance data for running HDFS replication jobs and for completed jobs.

To view the performance data for a running HDFS replication policy, perform the following steps:

Procedure

1. From Cloudera Manager, select Replication > Replication Policies.
2. Locate the row for the policy, select the policy, and click Actions. Select Show History.
3. Click Download CSV, and then choose one of the following options to view the performance report:
 - Performance. Click to download a CSV file which contains a summary report about the performance of the replication job. The performance summary report includes the last performance sample for each mapper that is working on the replication job.
 - Full Performance. Click to download the CSV file that contains the performance report of the job. The complete performance report includes all the samples taken for all mappers during the full execution of the replication job.

Replication Policies

Replication History

Name	test	Type	HDFS	Source	HDFS-1 (Cluster 1)	Destination	HDFS-1 (Cluster 1)	Next Run	None scheduled.
Start Time	Duration	Outcome	Files Expected	Files Copied	Files Failed	Files Deleted	Files Skipped		
✓ September 23, 2020 7:58 PM	1 min	Successful	80 (722.5 MiB)	17 (94.4 MiB)	0 (0 B)	0	63 (628.1 MiB)		
Started At	September 23, 2020 7:58 PM								
Duration	a few seconds								
Command Details	View								
MapReduce Job	job_1600880827337_0009								
HDFS Replication	<div>Download CSV ▾</div>								
Report	<div>Listing</div>								
Message	ed.								
> September 23, 2020 7:43		Successful	63 (628.1 MiB)	15 (93.2 MiB)	0 (0 B)	0	48 (534.9 MiB)		
> September 23, 2020 7:41		Successful	48 (534.9 MiB)	13 (92 MiB)	0 (0 B)	0	35 (442.9 MiB)		
> September 23, 2020 7:39		Successful	35 (442.9 MiB)	11 (90.8 MiB)	0 (0 B)	0	24 (352.2 MiB)		
> September 23, 2020 7:37		Successful	24 (352.2 MiB)	9 (89.6 MiB)	0 (0 B)	0	15 (262.6 MiB)		

4. To view the data, open the file in a spreadsheet program such as Microsoft Excel.

What to do next

The following table shows the columns that you can view in the CSV file:

Table 2: HDFS Performance Report Columns

Performance Data Columns	Description
Timestamp	Time when the performance data was collected.
Host	Name of the host where the YARN or MapReduce job was running.
Bytes Copied	Number of bytes copied for the file currently being copied.
Time Elapsed (ms)	Total time elapsed in milliseconds for the copy operation of the file currently being copied.
Files Copied	Number of files copied.
Avg Throughput (KB/s)	Average throughput since the start of the file currently being copied in kilobytes per second.
Last File (bytes)	File size of the last file in bytes.
Last File Time (ms)	Time taken to copy the last file in milliseconds.
Last file throughput (KB/s)	Throughput since the start of the last file being copied in kilobytes per second.

In addition to the performance reports, you can view the reports of files with errors, files that are deleted, and files that are skipped during the replication job. To view the reports, perform the following steps:

- On the Replication Policies page, locate the policy and click Actions > Show History.

The Replication History page for the replication policy appears. Expand to view the replication job details.

- Click Download CSV for the following options:
 - Listing - Click to download the CSV file that contains the replication report. The file lists the list of files and directories copied during the replication job.
 - Status - Click to download the CSV file that contains the complete status report. The file contains the full status report of the files where the status of the replication is one of the following:
 - ERROR – An error occurred and the file was not copied.
 - DELETED – A deleted file.
 - SKIPPED – A file where the replication was skipped because it was up-to-date.
 - Error Status Only - Click to download the CSV file that contains the status report of all copied files with errors. The file lists the status, path, and message for the copied files with errors.
 - Deleted Status Only - Click to download the CSV file that contains the status report of all deleted files. The file lists the status, path, and message for the databases and tables that were deleted.
 - Skipped Status Only - Click to download the CSV file that contains the status report of all skipped files. The file lists the status, path, and message for the databases and tables that were skipped.
 - Performance - Click to download a CSV file which contains a summary report about the performance of the running replication job. The performance summary report includes the last performance sample for each mapper that is working on the replication job.
 - Full Performance - Click to download the CSV file that contains the performance report of the job. The performance report shows the samples taken for all the mappers during the full execution of the replication job.

To view the data, open the file in a spreadsheet program such as Microsoft Excel.

The performance data is collected every two minutes. Therefore, no data is available during the initial execution of a replication job because not enough samples are available to estimate throughput and other reported data.

A sample CSV file, as presented in Excel, is shown here:

replication_hdfs_performance-217

Timestamp	Host	Bytes Copied	Time Elapsed (ms)	Files Copied	Avg Throughput (KB/s)	Last file (bytes)	Last file time (ms)	Last file throughput (KB/s)
09/23/2020 19:58:43.255	vkt-2.vkt.root.hwx.site	Bytes Copied: 390057	Time Elapsed (ms): 11	Files Copied: 1	Throughput (KB/s): 34628.64	Last file (bytes): 390057	Last time (ms): 11	Last throughput (KB/s): 34628.64
09/23/2020 19:58:44.470	vkt-3.vkt.root.hwx.site	Bytes Copied: 71693	Time Elapsed (ms): 6	Files Copied: 1	Throughput (KB/s): 11668.78	Last file (bytes): 71693	Last time (ms): 6	Last throughput (KB/s): 11668.78
09/23/2020 19:58:43.469	vkt-2.vkt.root.hwx.site	Bytes Copied: 860523	Time Elapsed (ms): 23	Files Copied: 1	Throughput (KB/s): 36537.15	Last file (bytes): 860523	Last time (ms): 23	Last throughput (KB/s): 36537.15
09/23/2020 19:58:43.556	vkt-2.vkt.root.hwx.site	Bytes Copied: 821908	Time Elapsed (ms): 17	Files Copied: 1	Throughput (KB/s): 47214.38	Last file (bytes): 821908	Last time (ms): 17	Last throughput (KB/s): 47214.38
09/23/2020 19:58:43.351	vkt-2.vkt.root.hwx.site	Bytes Copied: 636056	Time Elapsed (ms): 15	Files Copied: 1	Throughput (KB/s): 41409.90	Last file (bytes): 636056	Last time (ms): 15	Last throughput (KB/s): 41409.90
09/23/2020 19:58:43.418	vkt-2.vkt.root.hwx.site	Bytes Copied: 687113	Time Elapsed (ms): 15	Files Copied: 1	Throughput (KB/s): 44733.92	Last file (bytes): 687113	Last time (ms): 15	Last throughput (KB/s): 44733.92

Note the following limitations and known issues:

- If you click the CSV download too soon after the replication job starts, Cloudera Manager returns an empty file or a CSV file that has columns headers only and a message to try later when performance data has actually been collected.
- If you employ a proxy user with the form user@domain, performance data is not available through the links.
- If the replication job only replicates small files that can be transferred in less than a few minutes, no performance statistics are collected.
- For replication policies that specify the Dynamic Replication Strategy, statistics regarding the last file transferred by a MapReduce job hide previous transfers performed by that MapReduce job.
- Only the last trace per MapReduce job is reported in the CSV file.

Hive/Impala replication

Hive/Impala replication enables you to copy (replicate) your Hive metastore and data from one cluster to another and synchronize the Hive metastore and data set on the destination cluster with the source, based on a specified replication policy.

Minimum Required Role: [Replication Administrator](#) (also provided by Full Administrator)

The destination cluster must be managed by the Cloudera Manager Server where the replication is being set up, and the source cluster can be managed by that same server or by a peer Cloudera Manager Server.



Note: Because of the warehouse directory changes between CDH clusters and CDP Private Cloud Base, Hive replication does not copy the table data from the database and tables specified in the source cluster. But the replication job gets successfully run without any disruptions. While replicating from CDH clusters to CDP Private Cloud Base, it is recommended that the HDFS Destination Path is defined. If HDFS Destination Path is not defined and Replicate HDFS File is set as true, the data is replicated with the original source name. For example, the replicated table data was to reside under /warehouse/tablespace/external/hive directory but the data was replicated to /user/hive/warehouse location. Also, not defining HDFS Destination Path before the replication process can result in a large chunk of HDFS space being used for unwanted data movement.



Important: Since Hive3 has a different default table type and warehouse directory structure, the following changes apply while replicating Hive data from CDH5 or CDH6 versions to CDP Private Cloud Base:

- When you replicate from a CDH cluster to a CDP Private Cloud Base cluster, all tables become External tables during Hive replication. This is because the default table type is ACID in Hive3, which is the only managed table type. As of this release, Replication Manager does not support Hive2 -> Hive3 replication into ACID tables and all the tables will necessarily be replicated as External tables.
- Replicated tables will be created under external Hive warehouse directory set by hive.metastore.warehouse.external.dir Hive configuration parameter. Users have to make sure that this has a different value than hive.metastore.warehouse.dir Hive configuration parameter, that is the location of Managed tables.
- If users want to replicate the same database from Hive2 to Hive3 (that will have different paths by design), they need to use Force Overwrite option per policy to avoid any mismatch issues.



Note: While replicating from Sentry to Ranger, the minimum supported Cloudera Manager version is 6.3.1 and above.

Configuration notes:

- If the hadoop.proxyuser.hive.groups configuration has been changed to restrict access to the Hive Metastore Server to certain users or groups, the hdfs group or a group containing the hdfs user must also be included in the list of groups specified for Hive/Impala replication to work. This configuration can be specified either on the Hive service as an override, or in the core-site HDFS configuration. This applies to configuration settings on both the source and destination clusters.

- If you configured on the target cluster for the directory where HDFS data is copied during Hive/Impala replication, the permissions that were copied during replication, are overwritten by the HDFS ACL synchronization and are not preserved



Note: If your deployment includes tables backed by Kudu, Replication Manager filters out Kudu tables for a Hive replication in order to prevent data loss or corruption.

Host selection for Hive/Impala replication

If your cluster has Hive clients installed on hosts with limited resources, Hive/Impala replication may use these hosts to run commands for the replication, which can cause the performance of the replication to degrade.

About this task

To improve performance, you can specify the hosts (a "white list") to use during replication so that the lower-resource hosts are not used.

Procedure

1. Click Clusters Hive Configuration .
2. Type Hive Replication in the search box.
3. Locate the Hive Replication Environment Advanced Configuration Snippet (Safety Valve) property.
4. Add the HOST_WHITELIST property. Enter a comma-separated list of hostnames to use for Hive/Impala replication.

For example, HOST_WHITELIST=host-1.mycompany.com,host-2.mycompany.com.

5. Enter a Reason for change, and then click Save Changes to commit the changes.

Hive tables and DDL commands

The following applies when using the drop table and truncate table DDL commands.

- If you configure replication of a Hive table and then later drop that table, the table remains on the destination cluster. The table is not dropped when subsequent replications occur.
- If you drop a table on the destination cluster, and the table is still included in the replication job, the table is re-created on the destination during the replication.
- If you drop a table partition or index on the source cluster, the replication job also drops them on the destination cluster.
- If you truncate a table, and the Delete Policy for the replication job is set to Delete to Trash or Delete Permanently, the corresponding data files are deleted on the destination during a replication.

Replication of parameters

Parameters of databases, tables, partitions, and indexes are replicated by default during Hive/Impala replications.

You can disable replication of parameters:

1. Log in to the Cloudera Manager Admin Console.
2. Go to the Hive service.
3. Click the Configuration tab.
4. Search for "Hive Replication Environment Advanced Configuration Snippet"
5. Add the following parameter:

```
REPLICATE_PARAMETERS=false
```

6. Click Save Changes.

Hive replication in dynamic environments

To use Replication Manager for Hive replication in environments where the Hive Metastore changes, such as when a database or table gets created or deleted, additional configuration is needed.

Procedure

1. Open the Cloudera Manager Admin Console.
2. Search for the HDFS Client Advanced Configuration Snippet (Safety Valve) for hdfs-site.xml property on the source cluster.
3. Add the following properties:
 - a) Name: replication.hive.ignoreDatabaseNotFound
Value: true
 - b) Name: replication.hive.ignoreTableNotFound
Value: true
4. Save the changes.
5. Restart the HDFS service.

Creating a Hive/Impala replication policy

You must set up your clusters before you create a Hive/Impala replication policy.

Before you begin

The Apache Ranger access policy model consists of the following components:

- Specification of the resources that you can apply to a replication policy which includes the HDFS files and directories; Hive databases, tables, and columns; and HBase tables, column-families, and columns.
- Specification of access conditions for specific users and groups.

Replication Manager functions consistently across HDFS and Hive:

- Replication policies can be set up on files or directories in HDFS and on external tables in Hive—without manual translation of Hive datasets to HDFS datasets, or vice versa. Hive Metastore information is also replicated.
- Applications that depend on external table definitions stored in Hive, operate on both replica and source as table definitions are updated.
- Set the Ranger policy for hdfs user on target cluster to perform all operations on all databases and tables. The same user role is used to import Hive Metastore. The hdfs user should have access to all Hive datasets, including all operations. Otherwise, Hive import fails during the replication process. To provide access, perform the following steps:
 1. Log in to Ranger Admin UI.
 2. Navigate to the Service Manager Hadoop_SQL Policies Access section, and provide hdfs user permission to the all-database, table, column policy name.
- On the target cluster, the hive user must have Ranger admin privileges. The same hive user performs the metadata import operation.

Procedure

1. If the source cluster is managed by a different Cloudera Manager server than the destination cluster, configure a peer relationship.

2. Go to the Cloudera Manager Replication Policies page, click Create Replication Policy.

Replication Policies

Search

Filters

STATUS

- Failed 0
- Succeeded 0
- Running 0
- Disabled 0
- Dry-run 0

TYPE

- HDFS 0
- HDFS-S3 0
- Hive 0
- Hive-S3 0
- HDFS-ADLS 0
- Hive-ADLS 0

Actions for Selected

Create Replication Policy

No replication policies.

3. Select Hive External Table Replication Policy.

Replication Policies

Search

Filters

STATUS

- Failed 0
- Succeeded 0
- Running 0
- Disabled 0
- Dry-run 0

TYPE

- HDFS 0
- HDFS-S3 0
- Hive 0
- Hive-S3 0
- HDFS-ADLS 0
- Hive-ADLS 0

Actions for Selected

Create Replication Policy


HDFS Replication Policy

Hive Replication Policy

No replication policies.

4. In the General tab, configure the following options:

Option	Description
Name	Enter a unique name for the replication policy.
Source	Select the cluster with the Hive service you want to replicate.
Destination	Select the destination for the replication. If there is only one Hive service managed by Cloudera Manager available as a destination, this is specified as the destination. If more than one Hive service is managed by this Cloudera Manager, select from among them.
Use HDFS Destination	Select this option based on the type of destination cluster you plan to use.
Import Sentry permissions	Select one of the following permissions: <ul style="list-style-type: none"> Do not import Sentry Permissions (Default) If Sentry permissions were exported from the CDH cluster, import both Hive object and URL permissions If Sentry permissions were exported from the CDH cluster, import only Hive object permissions



Option	Description
Replicate All	<p>Select the option to replicate all the Hive databases from the source.</p> <p>To replicate only selected databases, clear the option and enter the database name(s) and tables you want to replicate.</p> <ul style="list-style-type: none"> Specify multiple databases and tables using the plus symbol to add more rows to the specification. Specify multiple databases on a single line by separating their names with the pipe () character. For example: mydbname1 mydbname2 mydbname3. Use regular expressions in the database or table fields as shown in the following examples: <ul style="list-style-type: none"> To specify any database or table name, enter the following regular expression: <pre>[\w] . +</pre> To specify any database or table except the one named 'myname', enter the following regular expression: <pre>(? !myname\b) . +</pre> To specify all the tables in the db1 and db2 databases, enter the following regular expression: <pre>db1 db2 [\w_] +</pre> To specify all the tables of the db1 and db2 databases (alternate method), enter the following regular expression: <pre>db1 [\w_] +</pre> <p>Click + icon and enter the following expression:</p> <pre>db2 [\w_] +</pre>
Run As Username	<p>Enter the username to run the MapReduce job. By default, MapReduce jobs run as hdfs. To run the MapReduce job as a different user, enter the user name. If you are using Kerberos, you must provide a user name here, and it must have an ID greater than 1000.</p> <p> Note: The user running the MapReduce job should have read and execute permissions on the Hive warehouse directory on the source cluster. If you configure the replication job to preserve permissions, superuser privileges are required on the destination cluster.</p>
Run on peer as Username	Enter the username if the peer cluster is configured with a different superuser. This is applicable in a kerberized environment.

5. In the Resources tab, configure the following options:



Option	Description
Scheduler Pool	<p>(Optional) Enter the name of a resource pool in the field. The value you enter is used by the MapReduce Service you specified when Cloudera Manager executes the MapReduce job for the replication. The job specifies the value using one of these properties:</p> <ul style="list-style-type: none"> MapReduce – Fair scheduler: mapred.fairscheduler.pool MapReduce – Capacity scheduler: queue.name YARN – mapreduce.job.queue.name

Option	Description
Maximum Map Slots	Enter the number of map slots per mapper, as required. The default value is 20.
Maximum Bandwidth	Enter the bandwidth per mapper, as required. The default is 100 MB.
Replication Strategy	<p>Choose Static or Dynamic. Determines whether the file replication tasks must be distributed among the mappers statically or dynamically. The default is Dynamic.</p> <p>Static replication distributes file replication tasks among the mappers up front to achieve a uniform distribution based on the file sizes. Dynamic replication distributes file replication tasks in small sets to the mappers, and as each mapper completes its tasks, it dynamically acquires and processes the next unallocated set of tasks.</p>

6. In the Advanced tab, you can specify an export location, modify the parameters of the MapReduce job that performs the replication, and set other options. You can select a MapReduce service (if there is more than one in your cluster) and change the following parameters:

Option	Description
Replicate HDFS Files	Clear the option to skip replicating the associated data files.
Replicate Impala Metadata	<p>If both the source and destination clusters use CDH 5.7.0 or later up to and including 5.11.x, select No to avoid redundant replication of Impala metadata. This option appears if both source and destination clusters support this functionality.</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> Yes – replicates the Impala metadata. No – does not replicate the Impala metadata. Auto – Cloudera Manager determines whether or not to replicate the Impala metadata based on the CDH version. <p>To replicate Impala UDFs when the version of CDH managed by Cloudera Manager is 5.7 or lower, see Replicate Impala and Hive User Defined Functions (UDFs) for information on when to select this option.</p>
Force Overwrite	<p>Select the option to overwrite data in the destination metastore if incompatible changes are detected. For example, if the destination metastore was modified, and a new partition was added to a table, this option forces deletion of that partition, overwriting the table with the version found on the source.</p> <p> Important: If the Force Overwrite option is not selected, and the Hive/Impala replication process detects incompatible changes on the source cluster, Hive/Impala replication fails. This sometimes occurs with recurring replications, where the metadata associated with an existing database or table on the source cluster changes over time.</p>
Export Path	<p>Specify a path to override the default HDFS location for the export file.</p> <p>By default, Hive metadata is exported to a default HDFS location (/user/\${**user.name**}/.cm/hive) and then imported from this HDFS file to the destination Hive metastore. In this example, user.name is the process user of the HDFS service on the destination cluster.</p> <p> Note: In a Kerberized cluster, the HDFS principal on the source cluster must have read, write, and execute access to the Export Path directory on the destination cluster.</p>

Option	Description
Number of concurrent HMS connections	<p>Enter the number of concurrent Hive Metastore connections. The connections are used to concurrently import and export metadata from Hive. Increase the number of threads to improve Replication Manager performance. By default, a new replication policy uses 5 connections.</p> <ul style="list-style-type: none"> a. If you set the value to 1 or more, Replication Manager uses multi-threading with the number of connections specified. b. If you set the value to 0 or fewer, Replication Manager uses single threading and a single connection. Note that the source and destination clusters must run a Cloudera Manager version that supports concurrent HMS connections, Cloudera Manager 5.15.0 or higher and Cloudera Manager 6.1.0 or higher.
HDFS Destination Path	<p>Enter a path to override the default path.</p> <p>By default, Hive HDFS data files (for example, /user/hive/warehouse/db1/t1) are replicated to a location relative to "/" (in this example, to /user/hive/warehouse/db1/t1).</p> <p>For example, if you enter /ReplicatedData, the data files are replicated to /ReplicatedData/user/hive/warehouse/db1/t1.</p>
MapReduce Service	Select the MapReduce or YARN service to use.
Log path	Enter an alternate path for the logs.
Description	Enter a description of the replication policy.

Option	Description
Error Handling	<p>Select the following options based on your requirements:</p> <ul style="list-style-type: none"> • Skip Checksum Checks - Determines whether to skip checksum checks on the copied files. If selected, checksums are not validated. Checksums are checked by default. <p> Important: You must skip checksum checks to prevent replication failure due to non-matching checksums in the following cases:</p> <ul style="list-style-type: none"> • Replications from an encrypted zone on the source cluster to an encrypted zone on a destination cluster. • Replications from an encryption zone on the source cluster to an unencrypted zone on the destination cluster. • Replications from an unencrypted zone on the source cluster to an encrypted zone on the destination cluster. <p>Checksums are used for two purposes:</p> <ul style="list-style-type: none"> • To skip replication of files that have already been copied. If Skip Checksum Checks is selected, the replication job skips copying a file if the file lengths and modification times are identical between the source and destination clusters. Otherwise, the job copies the file from the source to the destination. • To redundantly verify the integrity of data. However, checksums are not required to guarantee accurate transfers between clusters. HDFS data transfers are protected by checksums during transfer and storage hardware also uses checksums to ensure that data is accurately stored. These two mechanisms work together to validate the integrity of the copied data. <ul style="list-style-type: none"> • Skip Listing Checksum Checks - Determines whether to skip checksum check when comparing two files to determine whether they are same or not. If skipped, the file size and last modified time are used to determine if files are the same or not. Skipping the check improves performance during the mapper phase. Note that if you select the Skip Checksum Checks option, this check is also skipped. • Abort on Error - Determines whether to abort the job on an error. If selected, files copied up to that point remain on the destination, but no additional files are copied. Abort on Error is not selected by default. • Abort on Snapshot Diff Failures - If a snapshot diff fails during replication, Replication Manager uses a complete copy to replicate data. If you select this option, the Replication Manager aborts the replication when it encounters an error instead.
Preserve	<p>Determines whether to preserve the Block Size, Replication Count, and Permissions as they exist on the source file system, or to use the settings as configured on the destination file system. By default, settings are preserved on the source.</p> <p> Note: You must be running as a superuser to preserve permissions. Use the Run As Username option to ensure that is the case.</p>

Option	Description
Delete Policy	<p>Determines whether files that were deleted on the source should also be deleted from the destination directory. This policy also determines the handling of files in the destination location that are unrelated to the source. Options include:</p> <ul style="list-style-type: none"> Keep Deleted Files - Retains the destination files even when they no longer exist at the source. (This is the default.). Delete to Trash - If the HDFS trash is enabled, files are moved to the trash folder. Delete Permanently - Uses the least amount of space; use with caution. This option does not delete the files and directories in the top level directory. This is in line with rsync/Hadoop DistCp behavior.
Alerts	<p>Determines whether to generate alerts for various state changes in the replication workflow. You can alert on failure, on start, on success, or when the replication workflow is aborted.</p>

7. Click Save Policy.

What to do next



Note: If your replication job takes a long time to complete, and tables change before the replication finishes, the replication may fail. Consider making the Hive Warehouse Directory and the directories of any external tables snapshottable, so that the replication job creates snapshots of the directories before copying the files.

Sentry to Ranger replication for Hive replication policies

When you create or edit a Hive replication policy, you can choose to migrate the Sentry policies for Hive objects, Impala data, and URLs that are being replicated. The Replication Manager converts the Sentry policies to Ranger policies for the migrated data in the target cluster. The minimum supported Cloudera Manager version 6.3.1 and above is required to replicate Sentry policies to Ranger.

In a Hive replication policy, if you choose the If Sentry permissions were exported from the CDH cluster, import both Hive object and URL permissions or If Sentry permissions were exported from the CDH cluster, import only Hive object permissions option, the Replication Manager performs the following tasks automatically during the replication job run:

1. Exports each Sentry policy as a single JSON file using the authzmigrator tool. The JSON file contains a list of resources, such as URI, database, table, or column and the policies that apply to it.
2. Copies the exported Sentry policies to the target cluster using the DistCp tool.
3. Ingests the Sentry policies into Ranger after filtering the policies related to the replication job using the authzmigrator tool through the Ranger rest endpoint. To filter the policies, the Replication Manager uses a filter expression that is passed to the authzmigrator tool by Cloudera Manager.

If you are replicating a subset of the tables in a database, database-level policies get converted to equivalent table-level policies for each table being replicated. (For example, ALL on database -> ALL on table individually for each table replicated).

There will be no reference to the original role names in Ranger. The permissions are granted directly to groups and users with respect to the resource and not the role. This is a different format to the Sentry to Ranger migration during an in-place upgrade to CDP Private Cloud Base, which does import and use the Sentry roles.

Regardless of whether a policy was modified or not, each policy will be re-created on each replication. If you wish to continue scheduling data replication but you also want to modify the target cluster's Ranger policies (and keep those modifications), you should disable the Sentry to Ranger migration on subsequent runs by editing the replication policy and choose the Do not import Sentry Permissions (Default) option.

Replication of Impala and Hive User Defined Functions (UDFs)

By default, for clusters where the version of CDH is 5.7 or higher, Impala and Hive UDFs are persisted in the Hive Metastore and are replicated automatically as part of Hive/Impala replications.

To replicate Impala UDFs when the version of CDH managed by Cloudera Manager is 5.6 or lower, see for information on when to select the Replicate Impala Metadata option on the Advanced tab when creating a Hive/Impala replication policy.

After a replication has run, you can see the number of Impala and Hive UDFs that were replicated during the last run of the schedule on the Replication Policies page. You can also view the number of replicated UDFs on the Replication History page for previously-run replications.

Monitoring the performance of Hive/Impala replication policies

You can monitor the progress of a Hive/Impala replication policy using performance data that you download as a CSV file from the Cloudera Manager Admin console.

This file contains information about the tables and partitions being replicated, the average throughput, and other details that can help diagnose performance issues during Hive/Impala replications. You can view this performance data for running Hive/Impala replication jobs and for completed jobs.

To view the performance data for a running Hive/Impala replication:

1. From Cloudera Manager, select Replication > Replication Policies.
2. Locate the row for the policy, select the policy, and click Actions. Select Show History.
3. Click Download CSV for HDFS Replication Report, and then choose one of the following options to view the performance report:
 - Performance. Click to download a CSV file which contains a summary report about the performance of the replication job. The performance summary report includes the last performance sample for each mapper that is working on the replication job.
 - Full Performance. Click to download the CSV file that contains the performance report of the job. The complete performance report includes all the samples taken for all mappers during the full execution of the replication job.
4. To view the data, import the file into a spreadsheet program such as Microsoft Excel.

In addition to the performance reports, you can view the reports of files with errors, files that are deleted, and files that are skipped during the replication job. To view the reports, perform the following steps:

- On the Replication Policies page, locate the policy and click Actions > Show History.

The Replication History page for the replication policy appears. Expand to view the replication job details.

- Click Download CSV for the following options:
 - Listing - Click to download the CSV file that contains the replication report. The file lists the list of files and directories copied during the replication job.

- **Status** - Click to download the CSV file that contains the complete status report. The file contains the full status report of the files where the status of the replication is one of the following:
 - **ERROR** – An error occurred and the file was not copied.
 - **DELETED** – A deleted file.
 - **SKIPPED** – A file where the replication was skipped because it was up-to-date.
- **Error Status Only** - Click to download the CSV file that contains the status report of all copied files with errors. The file lists the status, path, and message for the copied files with errors.
- **Deleted Status Only** - Click to download the CSV file that contains the status report of all deleted files. The file lists the status, path, and message for the databases and tables that were deleted.
- **Skipped Status Only** - Click to download the CSV file that contains the status report of all skipped files. The file lists the status, path, and message for the databases and tables that were skipped.
- **Performance** - Click to download a CSV file which contains a summary report about the performance of the running replication job. The performance summary report includes the last performance sample for each mapper that is working on the replication job.
- **Full Performance** - Click to download the CSV file that contains the performance report of the job. The performance report shows the samples taken for all the mappers during the full execution of the replication job.

To view the data, open the file in a spreadsheet program such as Microsoft Excel.

The performance data is collected every two minutes. Therefore, no data is available during the initial execution of a replication job because not enough samples are available to estimate throughput and other reported data.

To view the performance data for a completed Hive/Impala replication schedule:

1. From Cloudera Manager, select Replication > Replication Policies.
2. Locate the row for the policy, select the policy, and click Actions. Select Show History.
3. To view performance of the Hive phase, click Download CSV next to the Hive Replication Report label and select one of the following options:
 - **Results** - Downloads a listing of replicated tables in a CSV file.
 - **Performance** - Downloads a performance report for the Hive replication in a CSV file.



Note: The option to download the HDFS replication reports might not appear if the HDFS phase of the replication skipped all the HDFS files because they have not changed, or if the Replicate HDFS Files option (located on the Advanced tab when creating Hive/Impala replication schedules) is not selected.

See [Table 3: Hive Performance Report Columns](#) on page 34 for a description of the data in the HDFS performance reports.

4. To view the data, open the file in a spreadsheet program such as Microsoft Excel.

The performance data is collected every two minutes. Therefore, no data is available during the initial execution of a replication job because not enough samples are available to estimate throughput and other reported data.

The data returned by the CSV files downloaded from the Cloudera Manager Admin console has the following structure:

Table 3: Hive Performance Report Columns

Hive Performance Data Columns	Description
Timestamp	Time when the performance data was collected
Host	Name of the host where the YARN or MapReduce job was running.
DbName	Name of the database.
TableName	Name of the table.
TotalElapsedTimeSecs	Number of seconds elapsed from the start of the copy operation.

Hive Performance Data Columns	Description
TotalTableCount	Total number of tables to be copied. The value of the column will be -1 for replications where Cloudera Manager cannot determine the number of tables being changed.
TotalPartitionCount	Total number of partitions to be copied. If the source cluster is running Cloudera Manager 5.9 or lower, this column contains a value of -1 because older releases do not report this information.
DbCount	Current number of databases copied.
DbErrorCount	Number of failed database copy operations.
TableCount	Total number of tables (for all databases) copied so far.
CurrentTableCount	Total number of tables copied for current database.
TableErrorCount	Total number of failed table copy operations.
PartitionCount	Total number of partitions copied so far (for all tables).
CurrPartitionCount	Total number of partitions copied for the current table.
PartitionSkippedCount	Number of partitions skipped because they were copied in the previous run of the replication job.
IndexCount	Total number of index files copied (for all databases).
CurrIndexCount	Total number of index files copied for the current database.
IndexSkippedCount	Number of Index files skipped because they were not altered. Due to a bug in Hive, this value is always zero.
HiveFunctionCount	Number of Hive functions copied.
ImpalaObjectCount	Number of Impala objects copied.

Note the following limitations and known issues:

- If you click the CSV download too soon after the replication job starts, Cloudera Manager returns an empty file or a CSV file that has columns headers only and a message to try later when performance data has actually been collected.
- If you employ a proxy user with the form user@domain, performance data is not available through the links.
- If the replication job only replicates small files that can be transferred in less than a few minutes, no performance statistics are collected.
- For replication policies that specify the Dynamic Replication Strategy, statistics regarding the last file transferred by a MapReduce job hide previous transfers performed by that MapReduce job.
- Only the last trace of each MapReduce job is reported in the CSV file.

Enabling, disabling, or deleting a replication policy

When you create a new replication policy, it is automatically enabled. If you disable a replication policy, it can be re-enabled at a later time.

About this task

Managing replication schedules.

Procedure

1. From Cloudera Manager, select **Replication Replication Policies**.

2. Select Actions for Selected drop-down and Enable | Disable | Delete as applicable.

To enable, disable, or delete multiple replication schedules, you can select those policies from the Replication Policies page and repeat step 2.

Replicating data to Impala clusters

Impala metadata is replicated as part of regular Hive/Impala replication operations.

Replicating Impala Metadata

Impala metadata replication is performed as a part of Hive replication. Impala replication is only supported between two CDH clusters. The Impala and Hive services must be running on both clusters.

To enable Impala metadata replication, perform the following tasks:

1. Schedule a Hive replication.
2. Confirm that the Replicate Impala Metadata option is set to Yes on the Advanced tab in the Create Hive Replication dialog.

When you set the Replicate Impala Metadata option to Yes, Impala UDFs (user-defined functions) will be available on the target cluster, just as on the source cluster. As part of replicating the UDFs, the binaries in which they are defined are also replicated.



Note: To run queries or execute DDL statements on tables that have been replicated to a destination cluster, you must run the Impala `INVALIDATE METADATA` statement on the destination cluster to prevent queries from failing.

Invalidating Impala Metadata

For Impala clusters that do not use LDAP authentication, you can configure Hive replication jobs to automatically invalidate Impala metadata after replication completes. If the clusters use Sentry, the Impala user should have permissions to run `INVALIDATE METADATA`.

The configuration causes the Hive/Impala replication job to run the Impala `INVALIDATE METADATA` statement per table on the destination cluster after completing the replication. The statement purges the metadata of the replicated tables and views within the destination cluster's Impala upon completion of replication, allowing other Impala clients at the destination to query these tables successfully with accurate results. However, this operation is potentially unsafe if DDL operations are being performed on any of the replicated tables or views while the replication is running. In general, directly modifying replicated data/metadata on the destination is not recommended. Ignoring this can lead to unexpected or incorrect behavior of applications and queries using these tables or views.



Note: If the source contains UDFs, you must run the `INVALIDATE METADATA` statement manually and without any tables specified even if you configure the automatic invalidation.

To configure the option, perform the following tasks:

1. Schedule a Hive replication.
2. On the Advanced tab, select the Invalidate Impala Metadata on Destination option.

Alternatively, you can run the `INVALIDATE METADATA` statement manually for replicated tables.

Enabling replication between clusters with Kerberos Authentication

To enable replication between clusters, additional setup steps are required to ensure that the source and destination clusters can communicate.

Minimum Required Role: Cluster Administrator (also provided by Full Administrator)



Important: Cloudera Replication Manager works with clusters in different Kerberos realms even without a Kerberos realm trust relationship. The Cloudera Manager configuration properties Trusted Kerberos Realms and Kerberos Trusted Realms are used for Cloudera Manager and CDH configuration, and are not related to Kerberos realm trust relationships.

If you are using standalone DistCp between clusters in different Kerberos realms, you must configure a realm trust.


Ports

When using Replication Manager with Kerberos authentication enabled, Replication Manager requires all the ports listed on the following page [Port Requirements for Replication Manager](#) on page 7.

Additionally, the port used for the Kerberos KDC Server and KRB5 services must be open to all hosts on the destination cluster. By default, this is port 88.

Considerations for realm names

If the source and destination clusters each use Kerberos for authentication, use one of the following configurations to prevent conflicts when running replication jobs.

- If the clusters do not use the same KDC (Kerberos Key Distribution Center), Cloudera recommends that you use different realm names for each cluster. Additionally, if you are replicating across clusters in two different realms, see the steps for [HDFS, Hive, and Impala replication](#) on page 37 to setup trust between those clusters.
- You can use the same realm name if the clusters use the same KDC or different KDCs that are part of a unified realm, for example where one KDC is the master and the other is a worker KDC.
-  **Note:** If you have multiple clusters that are used to segregate production and non-production environments, this configuration could result in principals that have equal permissions in both environments. Make sure that permissions are set appropriately for each type of environment.



Important: If the source and destination clusters are in the same realm but do not use the same KDC or the KDCs are not part of a unified realm, the replication job will fail.

HDFS, Hive, and Impala replication

Configuring source and destination clusters.

1. On the hosts in the destination cluster, ensure that the `krb5.conf` file (typically located at `/etc/krb5.conf`) on each host has the following information:
 - The KDC information for the source cluster's Kerberos realm. For example:

```
[realms]
SRC.EXAMPLE.COM = {
  kdc = kdc01.src.example.com:88
  admin_server = kdc01.example.com:749
  default_domain = src.example.com
}
DST.EXAMPLE.COM = {
  kdc = kdc01.dst.example.com:88
  admin_server = kdc01.dst.example.com:749
  default_domain = dst.example.com
}
```

- Realm mapping for the source cluster domain. You configure these mappings in the [domain_realm] section. For example:

```
[domain_realm]
.dst.example.com = DST.EXAMPLE.COM
dst.example.com = DST.EXAMPLE.COM
.src.example.com = SRC.EXAMPLE.COM
src.example.com = SRC.EXAMPLE.COM
```



Important: If you have a scenario where the hostname(s) are inconsistent, you must navigate to Cloudera Manager > Host > All Hosts > Ensure that all those hosts are covered in a similar manner as seen in domain_realm section.

2. On the destination cluster, use Cloudera Manager to add the realm of the source cluster to the Trusted Kerberos Realms configuration property:
 - a. Go to the HDFS service.
 - b. Click the Configuration tab.
 - c. In the search field type Trusted Kerberos to find the Trusted Kerberos Realms property.
 - d. Click the plus sign icon, and then enter the source cluster realm.
 - e. Enter a Reason for change, and then click Save Changes to commit the changes.
3. Go to AdministrationSettings.
4. In the search field, type domain name.
5. In the Domain Name(s) field, enter any domain or host names you want to map to the destination cluster KDC. Use the plus sign icon to add as many entries as you need. The entries in this property are used to generate the domain_realm section in krb5.conf.
6. If domain_realm is configured in the Advanced Configuration Snippet (Safety Valve) for remaining krb5.conf, remove the entries for it.
7. Enter a Reason for change, and then click Save Changes to commit the changes.

Kerberos connectivity test

As part of Test Connectivity, Cloudera Manager tests for properly configured Kerberos authentication on the source and destination clusters that run the replication. Test Connectivity runs automatically when you add a peer for replication, or you can manually initiate Test Connectivity from the Actions menu.

This feature is available when the source and destination clusters run Cloudera Manager 5.12 or later. You can disable the Kerberos connectivity test by setting feature_flag_test_kerberos_connectivity to false with the Cloudera Manager API: `api/<version>/cm/config`.

If the test detects any issues with the Kerberos configuration, Cloudera Manager provides resolution steps based on whether Cloudera Manager manages the Kerberos configuration file.

Cloudera Manager tests the following scenarios:

- Whether both clusters have Kerberos enabled or not.
- Replication is supported from unsecure cluster to secure cluster starting Cloudera Manager 6.1 and later.
- Replication is not supported if the source cluster uses Kerberos and target cluster is unsecure.
- Whether both clusters are in the same Kerberos realm. Clusters in the same realm must share the same KDC or the KDCs must be in a unified realm.
- Whether clusters are in different Kerberos realms. If the clusters are in different realms, the destination cluster must be configured according to the following criteria:
 - Destination HDFS services must have the correct Trusted Kerberos Realms setting.
 - The krb5.conf file has the correct domain_realm mapping on all the hosts.
 - The krb5.conf file has the correct realms information on all the hosts.
- Whether the local and peer KDC are running on an available port. This port must be open for all hosts in the cluster. The default port is 88.

After Cloudera Manager runs the tests, Cloudera Manager makes recommendations to resolve any Kerberos configuration issues.

Kerberos recommendations

If Cloudera Manager manages the Kerberos configuration file, Cloudera Manager configures Kerberos correctly for you and then provides the set of commands that you must manually run to finish configuring the clusters.

If Cloudera Manager does not manage the Kerberos configuration file, Cloudera Manager provides the manual steps required to correct the issue.

Copying data between a secure and an insecure cluster using DistCp and WebHDFS

You can use distcp and WebHDFS to copy data between a secure cluster and an insecure cluster.

About this task

When copying data, ensure that you run the distcp commands from the secure cluster.

Procedure

1. On the secure cluster, set `ipc.client.fallback-to-simple-auth-allowed` to `true` in `core-site.xml`.

```
<property>
  <name>ipc.client.fallback-to-simple-auth-allowed</name>
  <value>true</value>
</property>
```

Alternatively, you can also pass this as a parameter when you run the distcp command. If you want to do that, move onto step 2.

2. On the insecure cluster, add the secured cluster's realm name to the insecure cluster's configuration.
 - a) In the Cloudera Manager Admin Console for the insecure cluster, navigate to **Clusters <HDFS cluster>**.
 - b) On the Configuration tab, search for **Trusted Kerberos Realms** and add the secured cluster's realm name.



Note: This does not require Kerberos to be enabled but is a necessary step to allow the simple auth fallback to happen in the `hdfs://` protocol.

- c) Save the change.
3. Use commands such as the following only from the secure cluster side.

```
#This example uses the insecure cluster as the source and the secure cluster as the destination
distcp webhdfs://<insecure_namenode>:9870 webhdfs://<secure_namenode>:9871

#This example uses the secure cluster as the source and the insecure cluster as the destination
distcp webhdfs://<secure_namenode>:9871 webhdfs://<insecure_namenode>:9870
```

If TLS is enabled, replace `webhdfs` with `swebhdfs`.

If you did not configure `ipc.client.fallback-to-simple-auth-allowed` and want to pass it as a parameter, run commands such as the following from the secure cluster:

```
#This example uses the insecure cluster as the source and the secure cluster (with TLS enabled) as the destination cluster. swwebhdfs is used instead of webhdfs when TLS is enabled.
hadoop distcp -D ipc.client.fallback-to-simple-auth-allowed=true webhdfs://<insecure_namenode>:9870 swwebhdfs://<secure_namenode>:9871
```

```
#This example uses the secure cluster (with TLS enabled) as the source cluster and the insecure cluster as the destination. swebhdfs is used instead of webhdfs when TLS is enabled.
hadoop distcp -D ipc.client.fallback-to-simple-auth-allowed=true swebhdfs://<secure_namenode>:9871 webhdfs://<insecure_namenode>:9870
```

Kerberos setup guidelines for Distcp between secure clusters

There are specific guidelines to consider while setting up Kerberos on secure CDP clusters for successfully performing distcp between them.

The guidelines mentioned here are only applicable for the following example deployment:

- You have two clusters, each in a different Kerberos realm (SOURCE and DESTINATION in this example)
- You have data that needs to be copied from SOURCE to DESTINATION
- A Kerberos realm trust exists, either between SOURCE and DESTINATION (in either direction), or between both SOURCE and DESTINATION and a common third realm (such as an Active Directory domain).

If your environment matches the one described above, use the following table to configure Kerberos delegation tokens on your cluster so that you can successfully distcp across two secure clusters. Based on the direction of the trust between the SOURCE and DESTINATION clusters, you can use the `mapreduce.job.hdfs-servers.token-renewal.exclude` property to instruct ResourceManagers on either cluster to skip or perform delegation token renewal for NameNode hosts.



Note: You must use the `mapreduce.job.hdfs-servers.token-renewal.exclude` parameter if both clusters use the HDFS Transparent Encryption feature.

Environment Type		Kerberos Delegation Token Setting
SOURCE trusts DESTINATION	Distcp job runs on the DESTINATION cluster	You do not need to set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property.
	Distcp job runs on the SOURCE cluster	Set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property to a comma-separated list of the hostnames of the NameNodes of the DESTINATION cluster.
DESTINATION trusts SOURCE	Distcp job runs on the DESTINATION cluster	Set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property to a comma-separated list of the hostnames of the NameNodes of the SOURCE cluster.
	Distcp job runs on the SOURCE cluster	You do not need to set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property.
Both SOURCE and DESTINATION trust each other	Set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property to a comma-separated list of the hostnames of the NameNodes of the DESTINATION cluster.	

Environment Type	Kerberos Delegation Token Setting
Neither SOURCE nor DESTINATION trusts the other	<p>If a common realm is usable (such as Active Directory), set the <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property to a comma-separated list of hostnames of the NameNodes of the cluster that is not running the distcp job. For example, if you are running the job on the DESTINATION cluster:</p> <ol style="list-style-type: none"> 1. kinit on any DESTINATION YARN Gateway host using an AD account that can be used on both SOURCE and DESTINATION. 2. Run the distcp job as the hadoop user: <pre>\$ hadoop distcp -Ddfs.namenode.kerberos.principal.pattern=* \ -Dmapreduce.job.hdfs-servers.token-renewal.exclude=SOURCE-nn-host1,SOURCE-nn-host2 \ hdfs://source-nn-nameservice/source/path \ /destination/path</pre> <p>By default, the YARN ResourceManager renews tokens for applications. The <code>mapreduce.job.hdfs-servers.token-renewal.exclude</code> property instructs ResourceManagers on either cluster to skip delegation token renewal for NameNode hosts.</p>

Replication of encrypted data

HDFS supports encryption of data at rest (including data accessed through Hive). This topic describes how replication works within and between encryption zones and how to configure replication to avoid failures due to encryption.

Encrypting data in transit between clusters

A source directory and destination directory may or may not be in an encryption zone. If the destination directory is in an encryption zone, the data on the destination directory is encrypted. If the destination directory is not in an encryption zone, the data on that directory is not encrypted, even if the source directory is in an encryption zone. Encryption zones are not supported in CDH versions 5.1 or lower.

When you configure encryption zones, you also configure a Key Management Server (KMS) to manage encryption keys. During replication, Cloudera Manager uses TLS/SSL to encrypt the keys when they are transferred from the source cluster to the destination cluster. When a HDFS replication command that specifies an encrypted source directory runs, Cloudera Manager temporarily copies the encryption keys from the source cluster to the destination cluster, using TLS/SSL (if configured for the KMS) to encrypt the keys. Cloudera Manager then uses these keys to decrypt the encrypted files when they are received from the source cluster before writing the files to the destination cluster.



Important: When you configure HDFS replication, you must select the Skip Checksum check property to prevent replication failure in the following cases:

- Replications from an encrypted zone on the source cluster to an encrypted zone on a destination cluster.
- Replications from an encryption zone on the source cluster to an unencrypted zone on the destination cluster.
- Replications from an unencrypted zone on the source cluster to an encrypted zone on the destination cluster.

Even when the source and destination directories are both in encryption zones, the data is decrypted as it is read from the source cluster (using the key for the source encryption zone) and encrypted again when it is written to the destination cluster (using the key for the destination encryption zone). The data transmission is encrypted if you have configured encryption for HDFS data transfer.



Note: The decryption and encryption steps happen in the same process on the hosts where the MapReduce jobs that copy the data run. Therefore, data in plain text only exists within the memory of the Mapper task. If a KMS is in use on either the source or destination clusters, and you are using encrypted zones for either the source or destination directories, configure TLS/SSL for the KMS to prevent transferring the key to the mapper task as plain text.

During replication, data travels from the source cluster to the destination cluster using distcp. For clusters that use encryption zones, configure encryption of KMS key transfers between the source and destination using TLS/SSL. See [Configuring TLS/SSL for the KMS](#).

To configure encryption of data transmission between source and destination clusters:

- Enable TLS/SSL for HDFS clients on both the source and the destination clusters. For instructions, see [Configuring TLS/SSL for HDFS, YARN, and MapReduce](#). You may also need to configure trust between the SSL certificates on the source and destination.
- Enable TLS/SSL for the two peer Cloudera Manager Servers. See [Manually Configuring TLS Encryption for Cloudera Manager](#).
- Encrypt data transfer using HDFS data transfer encryption. See [Configuring Encrypted Transport for HDFS](#).

The following blog post provides additional information about encryption with HDFS: <https://blog.cloudera.com/blog/2013/03/how-to-set-up-a-hadoop-cluster-with-network-encryption/>.

Security considerations

The user you specify with the Run As field when scheduling a replication job requires full access to both the key and the data directories being replicated. This is not a recommended best practice for KMS management. If you change permissions in the KMS to enable this requirement, you could accidentally provide access for this user to data in other encryption zones using the same key. If a user is not specified in the Run As field, the replication runs as the default user, hdfs.

To access encrypted data, the user must be authorized on the KMS for the encryption zones they need to interact with. The user you specify with the Run As field when scheduling a replication must have this authorization. The key administrator must add ACLs to the KMS for that user to prevent authorization failure.

Key transfer using the KMS protocol from source to the client uses the REST protocol, which requires that you configure TLS/SSL for the KMS. When TLS/SSL is enabled, keys are not transferred over the network as plain text.

Snapshots

You can create HBase and HDFS snapshots using Cloudera Manager or by using the command-line.

- HBase snapshots allow you to create point-in-time backups of tables without making data copies, and with minimal impact on RegionServers. HBase snapshots are supported for clusters running CDH 4.2 or higher.
- HDFS snapshots allow you to create point-in-time backups of directories or the entire filesystem without actually cloning the data. They can improve data replication performance and prevent errors caused by changes to a source directory. These snapshots appear on the filesystem as read-only directories that can be accessed just like other ordinary directories.

Cloudera Manager snapshot policies

Cloudera Manager enables the creation of snapshot policies that define the directories or tables to be snapshotted, the intervals at which snapshots should be taken, and the number of snapshots that should be kept for each snapshot interval.

For example, you can create a policy that takes both daily and weekly snapshots, and specify that seven daily snapshots and five weekly snapshots should be maintained.

Minimum Required Role: [Replication Administrator](#) (also provided by Full Administrator)



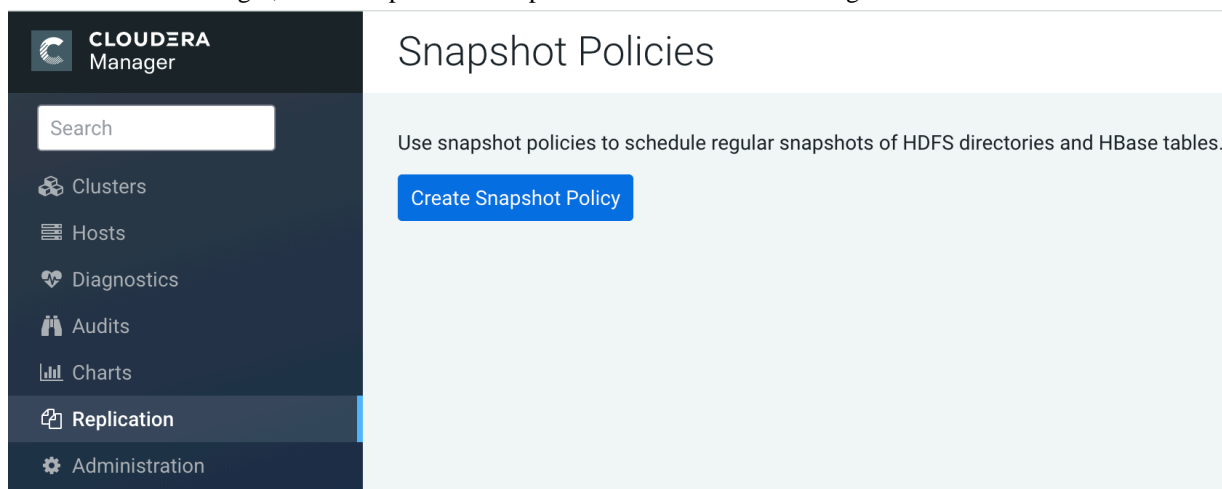
Note: You can improve the reliability of by also using snapshots.

Managing snapshot policies

You must enable an HDFS directory for snapshots to allow snapshot policies to be created for that directory.

To create a snapshot policy:

1. From Cloudera Manager, select **Replication Snapshot Policies** in the left navigation bar.



Existing snapshot policies are shown in a table.

2. To create a new policy, click **Create Snapshot Policy**.
3. From the drop-down list, select the service (HDFS or HBase) and cluster for which you want to create a policy.
4. Provide a name for the policy. Optionally, provide a description.
5. Specify the directories, namespaces or tables to include in the snapshot.



Important: Do not take snapshots of the root directory.

- For an HDFS service, select the paths of the directories to include in the snapshot. The drop-down list allows you to select only directories that are enabled for snapshotting. If no directories are enabled for snapshotting, a warning displays.

Click **+** to add a path and **=** to remove a path.

- For an HBase service, list the tables to include in your snapshot. You can use a [Java regular expression](#) to specify a set of tables. For example, `finance.*` matches all tables with names starting with `finance`. You can also create a snapshot for all tables in a given namespace, using the `{namespace}.*` syntax.
6. Specify the snapshot Schedule. You can schedule snapshots hourly, daily, weekly, monthly, or yearly, or any combination of those. Depending on the frequency you select, you can specify the time of day to take the snapshot, the day of the week, day of the month, or month of the year, and the number of snapshots to keep at each interval. Each time unit in the schedule information is shared with the time units of larger granularity. That is, the minute value is shared by all the selected schedules, hour by all the schedules for which hour is applicable, and so on. For example, if you specify that hourly snapshots are taken at the half hour, and daily snapshots taken at the hour 20, the daily snapshot will occur at 20:30.

To select an interval, check its box. Fields display where you can edit the time and number of snapshots to keep. For example:

7. Specify whether Alerts should be generated for various state changes in the snapshot workflow. You can alert on failure, on start, on success, or when the snapshot workflow is aborted.

8. Click Save Policy.

The new Policy displays on the Snapshot Policies page.

To edit or delete a snapshot policy:

1. From Cloudera Manager, select **Replication Snapshot Policies** in the left navigation bar.

Existing snapshot policies are shown in a table.

2. Click the Actions menu shown next to a policy and select **Edit** or **Delete**.

Snapshots history

The Snapshots History page displays information about Snapshot jobs that have been run or attempted.

The page displays a table of Snapshot jobs with the following columns:

Table 4: Snapshots History

Column	Description
Start Time	Time when the snapshot job started execution. Click to display details about the snapshot. For example: Click the View link to open the Managed scheduled snapshots Command page, which displays details and messages about each step in the execution of the command. For example:
Outcome	Displays whether the snapshot succeeded or failed.
Paths Tables Processed	HDFS snapshots: the number of Paths Processed for the snapshot. HBase snapshots: the number of Tables Processed for the snapshot.
Paths Tables Unprocessed	HDFS Snapshots: the number of Paths Unprocessed for the snapshot. HBase Snapshots: the number of Tables Unprocessed for the snapshot.
Snapshots Created	Number of snapshots created.
Snapshots Deleted	Number of snapshots deleted.
Errors During Creation	Displays a list of errors that occurred when creating the snapshot. Each error shows the related path and the error message.
Errors During Deletion	Displays a list of errors that occurred when deleting the snapshot. Each error shows the related path and the error message.

See [Table 4: Snapshots History](#) for more information about managing snapshots.

Orphaned snapshots

When you edit or delete a snapshot policy, the snapshots for the files, directories, or tables that were removed from the snapshot policy are retained. These are known as *orphaned* snapshots. These snapshots are not deleted automatically because they are no longer associated with a snapshot policy.

You can identify and delete these orphaned snapshots manually through Cloudera Manager, or by creating a command-line script that uses the HDFS or HBase snapshot commands.

To avoid orphaned snapshots, you can choose one of the following methods depending on your requirements.

- Delete the snapshots before you edit or delete the associated snapshot policy.

Cloudera Manager assigns the prefix `cm-auto` which is followed by a globally unique identifier (GUID) for every HDFS snapshot policy. You can view the snapshot prefix in the policy summary in the policy list, and in the delete modal window.



Note: Before you delete a snapshot policy, ensure that you record the snapshot names in the snapshot policy and the `cm-auto-guid` of the snapshot policy. This is because you cannot determine the snapshot names in the snapshot policy and the `cm-auto-guid` of the snapshot policy after you delete the snapshot policy, and the snapshot names also do not contain any recognizable references to its snapshot policy.

- Identify the orphaned snapshots for a deleted snapshot policy using its `cm-auto-guid`, and delete the snapshots.

Managing HDFS snapshots

This topic demonstrates how to manage HDFS snapshots using either Cloudera Manager or the command line.

For HDFS services, use the File Browser tab to view the HDFS directories associated with a service on your cluster. You can view the currently saved snapshots for your files, and delete or restore them. From the HDFS File Browser tab, you can:

- Designate HDFS directories to be "snapshottable" so snapshots can be created for those directories.
- Initiate immediate (unscheduled) snapshots of a HDFS directory.
- View the list of saved snapshots currently being maintained. These can include one-off immediate snapshots, as well as scheduled policy-based snapshots.
- Delete a saved snapshot.
- Restore an HDFS directory or file from a saved snapshot.
- Restore an HDFS directory or file from a saved snapshot to a new directory or file (Restore As).

Before using snapshots, note the following limitations:

- Snapshots that include encrypted directories cannot be restored outside of the zone within which they were created.
- The Cloudera Manager Admin Console cannot perform snapshot operations (such as create, restore, and delete) for HDFS paths with encryption-at-rest enabled. This limitation only affects the Cloudera Manager Admin Console and does not affect CDH command-line tools or actions not performed by the Admin Console, such as Replication Manager which uses command-line tools. For more information about snapshot operations, see [the Apache HDFS snapshots documentation](#).

Browsing HDFS directories

You can browse through the HDFS directories to select the right cluster.

To browse the HDFS directories to view snapshot activity:

1. From the Clusters tab, select your cluster HDFS service.
2. Go to the File Browser tab.

As you browse the directory structure of your HDFS, basic information about the directory you have selected is shown at the right (owner, group, and so on).

Enabling and disabling HDFS snapshots

For snapshots to be created, HDFS directories must be enabled for snapshots. You cannot specify a directory as part of a snapshot policy unless it has been enabled for snapshots.

Before you begin

Minimum Required Role: [Cluster Administrator](#) (also provided by Full Administrator).

Procedure

1. To enable an HDFS directory for snapshots, select your cluster HDFS service on the Clusters tab.
2. Go to the File Browser tab.
3. Go to the directory you want to enable for snapshots.
4. In the File Browser, click the drop-down menu next to the full file path and select Enable Snapshots.
5. To disable snapshots for a directory that has snapshots enabled, click Disable Snapshots from the drop-down menu button at the upper right. If snapshots of the directory exist, they must be deleted before snapshots can be disabled.



Note: Once you enable snapshots for a directory, you cannot enable snapshots on any of its subdirectories. Snapshots can be taken only on directories that have snapshots enabled.

Taking and deleting HDFS snapshots

To manage HDFS snapshots, first enable an HDFS directory for snapshots.

Minimum Required Role: [Replication Administrator](#) (also provided by Full Administrator)

Taking snapshots



Note: You can also schedule snapshots to occur regularly by creating a Snapshot policy.

1. From the Clusters tab, select your CDH HDFS service.
2. Go to the File Browser tab.
3. Go to the directory with the snapshot you want to restore.
4. Click the drop-down menu next to the full path name and select Take Snapshot.

The Take Snapshot screen displays.


5. Enter a name for the snapshot.
6. Click OK.

The Take Snapshot button is present, enabling an immediate snapshot of the directory.

7. To take a snapshot, click Take Snapshot, specify the name of the snapshot, and click Take Snapshot. The snapshot is added to the snapshot list.

Any snapshots that have been taken are listed by the time at which they were taken, along with their names and a menu button.

Deleting snapshots

1. From the Clusters tab, select your CDH HDFS service.
2. Go to the File Browser tab.
3. Go to the directory with the snapshot you want to delete.
4. In the list of snapshots, locate the snapshot you want to delete and click .
5. Select Delete.

Restoring Snapshots

Before you restore from a snapshot, ensure that there is adequate disk space.

1. From the Clusters tab, select your CDH HDFS service.
2. Go to the File Browser tab.
3. Go to the directory you want to restore.

4. In the File Browser, click the drop-down menu next to the full file path (to the right of the file browser listings) and select one of the following:
 - Restore Directory From Snapshot
 - Restore Directory From Snapshot As...

The Restore Snapshot screen displays.

5. Select Restore Directory From Snapshot As... if you want to restore the snapshot to a different directory. Enter the directory path to which the snapshot has to be restored. Ensure that there is enough space on HDFS to restore the files from the snapshot.



Note: If you enter an existing directory path in the Restore Directory From Snapshot As... field, the directory is overwritten.

6. Select one of the following:
 - Use HDFS 'copy' command - This option executes more slowly and does not require credentials in a secure cluster. It copies the contents of the snapshot as a subdirectory or as files within the target directory.
 - Use DistCp / MapReduce - This options executes more quickly and requires credentials (Run As) in secure clusters. It merges the target directory with the contents of the source snapshot. When you select this option, the following additional fields, which are similar to those available when configuring a replication, display under More Options:
 - When restoring HDFS data, if a MapReduce or YARN service is present in the cluster, DistributedCopy (distcp) is used to restore directories, increasing the speed of restoration. The Restore Snapshots screen HDFS (under More Options) allows selection of either MapReduce or YARN as the MapReduce service. For files, or if a MapReduce or YARN service is not present, a normal copy is performed.
 - Skip Checksum Checks - Whether to skip checksum checks (the default is to perform them). If checked, checksum validation will not be performed.

You must select the this property to prevent failure when restoring snapshots in the following cases:

- Restoring a snapshot within a single encryption zone.
- Restoring a snapshot from one encryption zone to a different encryption zone.
- Restoring a snapshot from an unencrypted zone to an encrypted zone.

Using snapshots with replication

Some replications, especially those that require a long time to finish, can fail because source files are modified during the replication process. You can prevent such failures by using Snapshots in conjunction with Replication. This use of snapshots is automatic with CDH versions 5.0 and higher. To take advantage of this, you must enable the relevant directories for snapshots (also called making the directory snapshottable).

When the replication job runs, it checks to see whether the specified source directory is snapshottable. Before replicating any files, the replication job creates point-in-time snapshots of these directories and uses them as the source for file copies. This ensures that the replicated data is consistent with the source data as of the start of the replication job. The latest snapshot for the subsequent runs is retained after the replication process is completed.

A directory is snapshottable because it has been enabled for snapshots, or because a parent directory is enabled for snapshots. Subdirectories of a snapshottable directory are included in the snapshot.


Hive/Impala replication using snapshots

Before you create Hive external table replication policies, ensure that you enable snapshots for the databases and directories that contain the required external tables. Before you replicate Impala tables, ensure that the storage locations for the tables and associated databases are also snapshottable.

For example, if the database resides in a custom location, such as /apps/folder1/folder2/[sales.db, marketing.db, hr.db, etc.], you can enable the snapshots at the following database or directory levels depending on your requirement:

- /apps/folder1/folder2/sales.db

- /apps/folder1/folder2/marketing.db
- /apps/folder1/folder2/hr.db

 **Note:** If you enable snapshots at the /apps, /apps/folder1, or /apps/folder1/folder2 level, large snapshots are created which might create performance and snapshot-related issues.

You can also isolate the database-level snapshots from each other so that the Hive external table replication policy replicates only the specified database.

The following table shows sample custom locations that contain the external tables and the recommended directory level to enable snapshots to isolate the database-level snapshots:

Sample custom location of external tables	Recommended directory level to enable snapshots
/data/folder1/folder2/sales/[table1, table2, table3 ... tablen]	/data/folder1/folder2/sales
/data/folder1/folder2/marketing/[table1, table2, table3 ... tablen]	/data/folder1/folder2/marketing
/data/folder1/folder2/hr/[table1, table2, table3 ... tablen]	/data/folder1/folder2/hr