

Replication policies in Replication Manager

Date published: 2022-11-18

Date modified: 2024-03-01



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 policies in Replication Manager.....	4
HDFS replication policies.....	4
Preparing clusters for HDFS replication policies.....	5
Creating an HDFS replication policy.....	6
Managing HDFS replication policy.....	9
Hive replication policies.....	12
Preparing clusters for Hive replication policies.....	12
Creating a Hive replication policy.....	14
Managing Hive replication policy.....	19
Ozone replication policies.....	22
Preparing clusters for Ozone replication policies.....	23
Configuring properties for OBS buckets to use in Ozone replication policies.....	24
Creating an Ozone replication policy.....	26
Managing Ozone replication policy.....	30
Appendix.....	32
Replication between clusters using Kerberos authentication.....	32
Configuring kerberized clusters for replication.....	33
Configuring user to replicate from unsecure to secure cluster.....	34
Kerberos connectivity test.....	34
Configuring SSL/TLS certificate exchange between two Cloudera Manager instances.....	35

Replication policies in Replication Manager

Replication Manager is a service in CDP Private Cloud Data Services. You can use this service to copy and migrate HDFS data, Hive external tables, and Ozone data between CDP Private Cloud Base 7.1.8 or higher clusters using Cloudera Manager version 7.7.3 or higher.

You can create HDFS replication policies, Hive replication policies, and Ozone replication policies to replicate HDFS data, Hive external tables, and Ozone data respectively. Before you create replication policies, you must verify whether the on-premises cluster versions are supported by Replication Manager, add the on-premises clusters in the Management Console, and configure the required ports for replication.

HDFS replication policies

You can use HDFS replication policies in CDP Private Cloud Data Services Replication Manager to copy or replicate HDFS files and directories between CDP Private Cloud Base 7.1.8 or higher clusters. Before you create HDFS replication policies, you must be aware of the guidelines related to the replication process and limitations related to HDFS replication policies.

Guidelines

To replicate HDFS data successfully using HDFS replication policies, you must ensure the following guidelines are followed during replication:

- All the source files in the directory are closed. This is because replication fails if the source files are open.



Tip: After replication completes, view the log for the replication job to identify the opened files.



Important: If you cannot ensure that all source files are closed, configure the **Additional Settings Abort on Error** option in the HDFS replication policy wizard to continue data replication despite errors.

- Source directory is not modified during replication. This is because the files that are added during replication do not get replicated, and if an existing file is deleted during replication, the replication fails.
- Log files are closed before the next replication job is initiated. This is because the log files are updated during replication.
- Maintain the latency between the source cluster NameNode and the destination cluster NameNode to less than 80 milliseconds for best performance. This is because of high latency among clusters might cause replication jobs to run slowly, but the job does not fail. You can test latency using the Linux `ping` command.

Limitations

- Maximum of 100 million files can be handled by a single replication job.
- Maximum of 10 million files can be handled by a replication policy that runs more frequently than once in 8 hours.
- Throughput of the replication job depends on the absolute read and write throughput of the source and destination clusters.



Tip: Perform regular rebalancing of HDFS clusters for efficient operation of replications.

Preparing clusters for HDFS replication policies

Before you create an HDFS replication policy, you must verify whether CDP Private Cloud Data Services Replication Manager supports the clusters for replication, ensure the required ports are open, and if the required CDP Private Cloud Base clusters are available as clusters in the Management Console. You can also enable Kerberos authentication for clusters supporting Kerberos.

About this task

Complete the following checklist to ensure that the clusters are ready to be used in an HDFS replication policy:

Procedure

- Have you added the required source and target clusters on the [Management Console Clusters](#) page?

For more information, see [Adding clusters to a CDP Private Cloud Data Services deployment](#).



Important:

Consider the following limitations before you add clusters:

- Only users with admin and poweruser roles can add clusters to use in Replication Manager.
- When Auto-TLS is enabled for a cluster, you must use the Cloudera Manager hostname to add a cluster.
- If a Cloudera Manager instance manages multiple clusters, you cannot include those clusters for replication in CDP Private Cloud Data Services Replication Manager.
- Are snapshots enabled in Cloudera Manager for the source cluster?



Tip: To enable snapshots for an HDFS directory, click Enable Snapshots for the required directory on the Cloudera Manager Clusters <HDFS Service> File Browser tab.

- Are the following ports open and available for Replication Manager?

Port	Service	Description
7180 or 7183	Cloudera Manager Admin Console HTTP	Open on the source cluster to enable source Cloudera Manager to communicate with the target Cloudera Manager.
80 or 443	Data transfer from secondary node for AWS / ADLS Gen2	Outgoing port. Open on all the HDFS nodes for AWS and ADLS Gen2.

- Do the clusters support Kerberos?

1. Enable Kerberos authentication and configure additional settings on the clusters to enable replication. Perform the following steps to accomplish this task:

- a. Enable Kerberos authentication. For more information, see [Enabling Kerberos authentication for CDP](#).
- b. Enable replication between clusters using Kerberos authentication. For more information, see [Configuring kerberized clusters for replication](#) on page 33.

2. Add the destination principal as a proxy user on the source cluster if you are using different Kerberos principals for the source and destination clusters.

For example, if you are using the `hdfssrc` principal on the source cluster and the `hdfsdest` principal on the destination cluster, add the following key-value pairs to the HDFS service Cluster-wide Advanced Configuration Snippet (Safety Valve) for `core-site.xml` property on the source cluster and then restart the HDFS service:

- Name = `hadoop.proxyuser.hdfsdest.groups`; Value = *
- Name = `hadoop.proxyuser.hdfsdest.hosts`; Value = *

- Do you need to replicate data securely? If so, ensure that the SSL/TLS certificate exchange between two Cloudera Manager instances that manage source and target clusters respectively is configured. For more information, see [Configuring SSL/TLS certificate exchange between two Cloudera Manager instances](#) on page 35.

Cloudera Manager provides data analytics that you can download and use to diagnose HDFS replication performance.

Creating an HDFS replication policy

You can create an HDFS replication policy in CDP Private Cloud Data Services Replication Manager to replicate HDFS data between CDP Private Cloud Base 7.1.8 or higher clusters.

Procedure

1. Log into CDP Private Cloud Data Services, and click Replication Manager.
2. Click Replication Policies.
3. On the **Replication Policies** page, click Create Policy.
4. On the **General** page, choose or enter the following information:

Option	Action to take
HDFS	Choose to create an HDFS replication policy.
Policy Name	Enter a unique name for the replication policy.
Description	Optional. Enter a brief description about the replication policy.

5. Click Next.
6. On the **Select Source** page, choose or enter the following information:





Option	Action to take
Source cluster	Choose the source cluster.
Source Path	Choose one of the following methods to determine the directory where source data resides on the source cluster: <ul style="list-style-type: none"> • Enter the complete directory path. • Click File Browser to view and navigate the existing directory list on the selected cluster. Select the required directory that you want to replicate.
Run As Username (on source)	Optional. The replication policy uses the <i>Default</i> username to replicate HDFS data. If you are using a kerberized cluster, enter the required username. The replication policy uses this username to replicate the data in the kerberized cluster.

7. Click Next.
8. On the **Select Destination** page, choose or enter the following information:

Option	Action to take
Destination cluster	Choose the target cluster.
Destination Path	Enter the directory path on the target cluster to which the replication policy replicates HDFS data.
Run as Username	Optional. The replication policy uses the <i>Default</i> username to replicate HDFS data. Enter another username if you want the replication policy to use it to replicate data.


9. Click Validate Policy.
Replication Manager verifies whether the details provided are correct.
10. Click Next.


11. On the *Schedule* page, choose or enter the following information:

Option	Action to take
Run Now	Choose to initiate data replication after the replication policy creation is complete. Choose the frequency to replicate data periodically, if required.
Schedule Run	<p>Choose to run the replication policy to replicate data at a later time. Choose the date and time for the first run, and then choose the frequency to replicate data periodically.</p> <p>Replication Manager ensures that the same number of seconds elapse between the runs. For example, if you choose the Start Time as January 19, 2022 11.06 AM and Interval as 1 day, Replication Manager runs the replication policy for the first time at the specified time in the timezone the replication policy was created in, and then runs it exactly after 1 day that is, after 24 hours or 86400 seconds.</p> <p> Tip: On the Replication Policies page, click   to change the timezone.</p>
Frequency	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> Does Not Repeat. Custom - In the Custom Recurrence dialog box, choose the time, date, and the frequency to run the policy. <p> Note: Ensure that the frequency in a schedule enables a job to finish before the next job starts. Also, ensure that the jobs based on the same policy do not overlap. If a job is not completed before another job starts, the second job does not run and the job status appears as Skipped. If a job is consistently skipped, you might need to modify the frequency of the job.</p>

12. Click Next.**13. On the *Additional Settings* page, enter the values as necessary for the attributes:**

Option	Description
YARN Queue Name	Enter the name of the YARN queue for the cluster to which the replication job is submitted if you are using Capacity Scheduler queues to limit resource consumption. The default value for this field is <i>default</i> .
Maximum Maps Slots	Set the maximum number of map tasks (simultaneous copies) per replication job. The default value is 20.
Maximum Bandwidth	<p>Adjust this setting so that each map task is throttled to consume only the specified bandwidth. The default value for the bandwidth is 100 MB per second for each mapper.</p> <p>Each map task (simultaneous copy) is restricted to consume only the specified bandwidth. This is not always exact. The map throttles back its bandwidth consumption during a copy in such a way that the net used bandwidth tends towards the specified value. You can adjust this setting so that each map task is throttled to consume only the specified bandwidth so that the net used bandwidth tends towards the specified value.</p>
Path Exclusion	Enter one or more regular expressions separated by comma. Replication Manager does not copy the subdirectories or files from the source that matches one of the specified regular expressions, to the target cluster.

Option	Description
Replication Strategy	<p>Choose one of the following replication strategies:</p> <ul style="list-style-type: none"> Static distributes file replication tasks among the mappers up front to achieve an uniform distribution based on the file sizes. Dynamic distributes the file replication tasks in small sets to the mappers, and as each mapper completes its tasks, it dynamically acquires and processes the next set of unallocated tasks. <p>The default replication strategy is Dynamic.</p>
MapReduce Service	Choose the MapReduce or YARN service.
Log Path	Enter an alternate path for the logs, if required.
Error Handling	<p>Select the following options as necessary:</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> Note: 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. 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 not selected by default. Abort on Snapshot Diff Failures - If a snapshot diff fails during replication, the replication policy uses a complete copy to replicate data. If you select this option, the policy aborts the replication when it encounters an error instead.

Option	Description
Preserve	<p>Choose the required options 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.</p> <p>When Permission is selected, and both the source and destination clusters support ACLs, replication preserves ACLs. Otherwise, ACLs are not replicated. When Extended attributes is selected, 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.)</p> <p>If you select one or more of the Preserve options and you are replicating to S3 or ADLS, the values of all of these items are saved in metadata files on S3 or ADLS. When you replicate from S3 or ADLS to HDFS, you can select which of these options you want to preserve.</p> <p> Note: To preserve permissions to HDFS, you must be running as a superuser on the destination cluster. Use the Run As Username option to set the username.</p>
Delete Policy	<p>Choose the required options to determine whether the 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 option. • Delete to Trash - If the HDFS trash is enabled, files are moved to the trash folder. This is not supported when replicating to S3 or ADLS. • Delete Permanently - Uses the least amount of space; use with caution.
Alerts	<p>Choose to generate alerts for various state changes in the replication workflow. You can choose to generate an alert On Failure, On Start, On Success, or On Abort of the replication job.</p> <p>You can configure alerts to be delivered by email or sent as SNMP traps. If alerts are enabled for events, you can search for and view the alerts on the Events tab, even if you do not have email notification configured. For example, if you choose Command Result that contains the Failed filter on the Diagnostics Events page, the alerts related to the On Failure alert for all the replication policies for which you have set the alert appear. For more information, see Managing Alerts and Configuring Alert Delivery.</p>

14. Click Create.

If you selected Run Now on the **Schedule** page, the replication job starts data replication after you click Create.

Results

The replication task appears on the **Replication Policies** page.

If the 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 running the replication jobs. For more information, see [Guidelines to consider for HDFS replication policies](#).

Managing HDFS replication policy

After you create an HDFS replication policy in CDP Private Cloud Data Services Replication Manager, you can perform and monitor various tasks related to the replication policy. You can view the job progress and replication




logs. You can edit the advanced options to optimize a job run. You can suspend a job and also activate a suspended job. You can edit the replication policy as necessary.

About this task

On the **Replication Policies** page, you can perform the following actions and tasks on a replication policy and its jobs:

Procedure

- When you click Actions for an HDFS replication policy, the following actions appear:

Action	Description
Edit*	<p>Change the replication policy options as required for non-expired policies that are in active or suspended state. Based on the schedule you choose, the replication policy replicates data.</p> <p>You can edit the replication policies to better align with changing requirements. For example, you might want to change the frequency of a policy depending on the data size and importance of the data being replicated.</p> <p> Note: A replication policy is associated with a cluster or a cluster pair, therefore you cannot change the clusters in the policy.</p> <p>Optionally, expand a replication policy on the Replication Policies page to edit the replication policy options which include frequency (start time cannot be modified if the policy has already started), queue name, maximum bandwidth, and maximum map slots.</p> <p> Tip: To optimize the replication policy performance, you can configure the queue name, maximum bandwidth, and maximum map slots as necessary.</p>
Rerun	Runs the replication policy.
Delete	Deletes the replication policy permanently.
Suspend	Suspends a running replication policy. Activate the replication policy, if required.
View Log	<p>Download, copy, or open the log to track the job and to troubleshoot any issues.</p> <p> Tip: On the Overview Issues & Updates panel, the Job Status column shows the current job status. If the job failed, click Failed to view the log details about the job.</p>
Collect diagnostic bundle	Generates a diagnostic bundle for the replication policy. You can download the bundle as a ZIP file to your machine.
<p>To view and use the replication policies with an empty name in Replication Manager, you must understand the following implementation:</p> <ul style="list-style-type: none"> If the Cloudera Manager API version is lower than 51, an existing replication policy with an empty name can be used and updated. However, if you edit the replication policy and provide a name for the replication policy in versions higher or equal to 51, you must ensure that the name conforms to the validation rules. If the Cloudera Manager API version is higher or equal to 51, it is mandatory that you provide a non-empty unique name to the replication policy to continue using it. This is because API version 51 and higher enforces the validation rules on all the replication policies. <p>To pass the replication policy name validation, you must ensure that the replication policy name is unique. The name can contain letters, numbers, and the _ / - characters. You must also ensure that it does not contain the characters % . ; \ nor any character that is not ASCII printable, which includes the ASCII characters less than 32 and the ASCII characters that are greater than or equal to 127.</p>	

- When you expand the policy details, the **Job History** panel appears.

You can view the following details on the panel:

- Previous jobs, current job, and one future scheduled job if any.
- Job details which include:

Job details	Description
Started	Timestamp when the job started.
Ended	Timestamp when the job ended.
Duration	Time taken to complete the job.
Progress	Current status of a running job.
Expected	Remaining number of files and bytes expected to be copied for a running job.
Copied	Number of files and bytes copied for a running job and completed job.
Failed	Number of files and bytes that failed to be copied for a completed job.
Deleted	Number of files deleted for a completed job.
Skipped	Remaining number of files and bytes skipped from copying for a running job and complete job.

- Click Actions to:

- Abort the job.
- Re-run an aborted or failed job.
- View Log for the job. You can download, copy, or open it to track the job and to troubleshoot any issues for the job.

- When you click a job on the **Job History** panel, the following tabs appear:

Tab	Description
General	Shows the following job details: <ul style="list-style-type: none"> Started at timestamp Duration to complete the job HDFS Replication Report to download the job statistics in CSV format Job status Message
Command Details	Shows the steps that Replication Manager ran for the job along with the timestamp.

- You can download the following CSV reports from the **General HDFS Replication Report** field to track the replication jobs and to troubleshoot issues:

Report	Description
Listing	Lists all the files and directories copied during the replication job.
Status	Shows the complete status report of each file as: <ul style="list-style-type: none"> an Error occurred and the file was not copied. a Deleted file. an up-to-date file for which the replication was Skipped.
Error Status	Status report of all the copied files with errors. Each file shows the status, path, and message for the copied files with errors.
Skipped Status	Status report of all skipped files. Each file lists the status, path, and message for the databases and tables that were skipped.
Deleted Status	Status report of all deleted files. Each file lists the status, path, and message for the databases and tables that were deleted.

Report	Description
Performance	Summary report about the performance of the running replication job which includes the last performance sample for each mapper that is working on the replication job.
Full Performance	Performance report of the job which includes the samples taken for all mappers during the replication job.



Note: The reports are generated based on the source Cloudera Manager response. If the Cloudera Manager response is interrupted or is not handled as expected, corresponding error messages appear in HTML format in the reports.

Hive replication policies

You can use Hive replication policies in CDP Private Cloud Data Services Replication Manager to replicate Hive metastore and Hive external tables between CDP Private Cloud Base 7.1.8 or higher clusters.

Before you replicate using Hive replication policies, consider the following limitations and guidelines:

- Managed table replication is not supported. Replication Manager translates the managed table from the source cluster to the target cluster as an external table. Replication Manager stores the replicated table as an external table.
- If the `hadoop.proxyuser.hive.groups` configuration is changed to restrict access to the Hive Metastore Server for certain users or groups, you must include the `hdfs` group or a group containing the `hdfs` user 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.xml` file in the HDFS configuration on both the source and destination clusters.
- Before you use the drop table and truncate table DDL commands, consider the following:
 - If you create a Hive replication policy for a Hive table and you drop the table in the source cluster after replication, 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 it on the destination cluster.
 - If you truncate a table, and the Delete Policy option during replication policy creation is set to Delete to Trash or Delete Permanently, the corresponding data files are deleted on the destination during subsequent replication jobs.

Preparing clusters for Hive replication policies

Before you create a Hive replication policy, you must verify whether CDP Private Cloud Data Services Replication Manager supports the clusters for replication, the required ports are open, and the clusters for replication are added in the Management Console. You can also enable Kerberos authentication for clusters supporting Kerberos.

About this task

Complete the following checklist to ensure that the clusters are ready to be used in a Hive replication policy:

Procedure

- Have you configured the all-database, table, column Ranger policy for the *hdfs* user on the target cluster to perform all the operations on all databases and tables?

The *hdfs* user role is used to import Hive Metastore and must have access to all Hive datasets, including all operations. Otherwise, Hive import fails during the replication process. On the target cluster, the *hive* user must have Ranger admin privileges. The same *hive* user performs the metadata import operation.



Tip: To provide access, go to the Ranger Admin UI Service Manager Hadoop_SQL Policies Access section, and provide *hdfs* user permission to the all-database, table, column policy name.

- Have you added the source and target clusters on the Management Console Clusters page?

For more information, see [Adding clusters to a CDP Private Cloud Data Services deployment](#).



Important: Consider the following limitations before you add clusters:

- Only users with admin and poweruser roles can add clusters to use in Replication Manager.
 - When Auto-TLS is enabled for a cluster, you must use the Cloudera Manager hostname to add a cluster.
 - If a Cloudera Manager instance manages multiple clusters, you cannot include those clusters for replication in CDP Private Cloud Data Services Replication Manager.
- Is the Hive Warehouse Directory enabled for snapshots to use in Hive replication policies?



Tip: Perform the following steps to accomplish this task:

- Search for `hive.metastore.warehouse.dir` on the Cloudera Manager *Hive service* Configuration tab.
- Enter the Hive warehouse directory path that is located in the HDFS file system for the Hive Warehouse Directory property. By default, the Hive warehouse directory is located in the `/user/hive/warehouse` location.

- Are the directories hosting the external tables that are not stored in the Hive warehouse directory snapshottable?

A directory is *snapshottable* if 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.



Tip: To enable snapshots for an HDFS directory, click Enable Snapshots for the required directory on the Cloudera Manager Clusters *HDFS Service* File Browser tab.

- Are the following ports open and available for Replication Manager?

Port	Service	Description
7180 or 7183	Cloudera Manager Admin Console HTTP	Open on the source cluster to enable source Cloudera Manager to communicate with the target Cloudera Manager.
80 or 443	Data transfer from secondary node for AWS / ADLS Gen2	Outgoing port. Open on all the HDFS nodes for AWS and ADLS Gen2.

- Are the following additional configurations configured? You can configure one or all of the configurations as necessary.
 - If your cluster has Hive clients installed on hosts with limited resources and the replication policies use these hosts to run the commands, the replication performance might degrade. To improve performance, you can specify the hosts (an “allowlist”) that can be used during the replication process so that the lower-resource hosts are not used. To specify the hosts, perform the following steps:
 - Go to the Cloudera Manager Clusters *Hive service* Configuration page.
 - Locate the Hive Replication Environment Advanced Configuration Snippet (Safety Valve) property.
 - Add the `HOST_WHITELIST` property, and 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`.
 - Click Save Changes.
 - Restart the Hive service.

- b) If the Hive Metastore is constantly updated, by activities such as creating or deleting databases or tables, you might want to configure the following properties to optimize replication performance:
 - a. Go to the Cloudera Manager Clusters *Hive service* Configuration page for the source cluster.
 - b. Search for the HDFS Client Advanced Configuration Snippet (Safety Valve) for `hdfs-site.xml` property.
 - c. Add the following key-value pairs:
 - Name = `replication.hive.ignoreDatabaseNotFound`; Value = `true`
 - Name = `replication.hive.ignoreTableNotFound`; Value = `true`
 - d. Click Save Changes.
 - e. Restart the HDFS service.
- c) Hive replication policies replicate parameters of databases, tables, partitions, and indexes by default. Perform the following steps to disable replication of parameters:
 - a. Go to the Cloudera Manager Clusters *Hive service* Configuration page for the source cluster.
 - b. Search for Hive Replication Environment Advanced Configuration Snippet property.
 - c. Add the following key-value pair:
 - Name = `REPLICATE_PARAMETERS`; Value = `false`
 - d. Click Save Changes.
 - e. Restart the Hive service.
- Do the clusters support Kerberos?
 1. Enable Kerberos authentication and configure additional settings on the clusters to enable replication. Perform the following steps to accomplish this task:
 - a. Enable Kerberos authentication. For more information, see [Enabling Kerberos authentication for CDP](#).
 - b. Enable replication between clusters using Kerberos authentication. For more information, see [Configuring kerberized clusters for replication](#) on page 33.
 2. Add the destination principal as a proxy user on the source cluster if you are using different Kerberos principals for the source and destination clusters.

For example, if you are using the `hdfssrc` principal on the source cluster and the `hdfsdest` principal on the destination cluster, add the following key-value pairs to the HDFS service Cluster-wide Advanced Configuration Snippet (Safety Valve) for `core-site.xml` property on the source cluster and then restart the HDFS service:

 - Name = `hadoop.proxyuser.hdfsdest.groups`; Value = `*`
 - Name = `hadoop.proxyuser.hdfsdest.hosts`; Value = `*`
- Do you need to replicate data securely? If so, ensure that the SSL/TLS certificate exchange between two Cloudera Manager instances that manage source and target clusters respectively is configured. For more information, see [Configuring SSL/TLS certificate exchange between two Cloudera Manager instances](#) on page 35.

Creating a Hive replication policy

You can create a Hive replication policy in CDP Private Cloud Data Services Replication Manager to replicate Hive external tables between CDP Private Cloud Base 7.1.8 or higher clusters.

Procedure


1. Log into CDP Private Cloud Data Services, and click Replication Manager.
2. Click Replication Policies.
3. On the **Replication Policies** page, click Create Policy.
4. On the **General** page, choose or enter the following information:

Option	Action to take
Hive	Choose to create a Hive replication policy.

Option	Action to take
Policy Name	Enter a unique name for the replication policy.
Description	Optional. Enter a brief description about the replication policy.


5. Click Next.

6. On the **Select Source** page, choose or enter the following information:

Option	Action to take
Source cluster	Choose the source cluster.
Source Databases and Tables	<p>Select one of the following options to choose the databases and tables to replicate:</p> <ul style="list-style-type: none"> Choose All Databases or enter * in the Select Databases and Tables field. Enter the database name, and then enter the table names separated by a whitespace. To enter multiple databases on a single line, separate the database names with the pipe () character. For example: <code>[***mydbname1***]/[***mydbname2***]/[***mydbname3***]</code>. Enter a regular expression (regex) to specify the databases and tables to be replicated. Use one of the following formats to specify databases and table names: <ul style="list-style-type: none"> <code>[w].+</code> to specify a database or table name. <code>(?![***excludename***]/b).+</code> to specify a database or table except the one named <i>excludename</i>. <code>[***db1***]/[***db2***] [w_]+</code> specifies all the tables in <i>db1</i> and <i>db2</i> databases.
Run As Username (on source)	<p>Optional. The replication policy uses the <i>Default</i> username to replicate data.</p> <p>Enter another username if you want the replication policy to use it to replicate data. Ensure that the user has the necessary permissions to replicate data.</p> <p> Note: You must provide a username for Kerberized clusters.</p>

7. Click Next.





8. On the **Select Destination** page, choose or enter the following information:

Option	Action to take
Destination cluster	Choose the target cluster.
Destination Path (Optional)	<p>Enter the path on the target cluster to which the policy replicates data.</p> <p>By default, Hive metadata is exported to a default HDFS location <code>/user/\${***user.name***}/cm/hive</code> and then imported from this HDFS file to the destination Hive metastore. In this example, <i>user.name</i> is the process user of the HDFS service on the destination cluster. To override the default HDFS location for this export file, you can enter specify a path in the Destination Path field.</p>
Run as Username	<p>Enter the user name to run the MapReduce job. By default, MapReduce jobs run as <i>hdfs</i>. 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 must have read and execute permissions on the Hive warehouse directory on the <i>source</i> cluster. If you configure the replication job to preserve permissions, superuser privileges are required on the <i>destination</i> cluster.</p>

9. Click Validate Policy.




Replication Manager verifies whether the details provided are correct.


10. Click Next.**11. On the **Schedule** page, choose or enter the following information:**


Option	Action to take
Run Now	Choose to initiate data replication after the replication policy creation is complete. Choose the frequency to replicate data periodically, if required.
Schedule Run	<p>Choose to run the replication policy to replicate data at a later time. Choose the date and time for the first run, and then choose the frequency to replicate data periodically.</p> <p>Replication Manager ensures that the same number of seconds elapse between the runs. For example, if you choose the Start Time as January 19, 2022 11.06 AM and Interval as 1 day, Replication Manager runs the replication policy for the first time at the specified time in the timezone the replication policy was created in, and then runs it exactly after 1 day that is, after 24 hours or 86400 seconds.</p> <p> Tip: On the Replication Policies page, click   to change the timezone.</p>
Frequency	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> Does Not Repeat. Custom - In the Custom Recurrence dialog box, choose the time, date, and the frequency to run the policy. <p> Note: Ensure that the frequency in a schedule enables a job to finish before the next job starts. Also, ensure that the jobs based on the same policy do not overlap. If a job is not completed before another job starts, the second job does not run and the job status appears as Skipped. If a job is consistently skipped, you might need to modify the frequency of the job.</p>

12. Click Next.**13. On the **Additional Settings** page, enter the values as necessary for the attributes:**

Option	Description
YARN Queue Name	Enter the name of the YARN queue for the cluster to which the replication job is submitted if you are using Capacity Scheduler queues to limit resource consumption. The default value for this field is default.
Maximum Maps Slots	Set the maximum number of map tasks (simultaneous copies) per replication job. The default value is 20.
Maximum Bandwidth	<p>Adjust this setting so that each map task is throttled to consume only the specified bandwidth. The default value for the bandwidth is 100 MB per second for each mapper.</p> <p>Each map task (simultaneous copy) is restricted to consume only the specified bandwidth. This is not always exact. The map throttles back its bandwidth consumption during a copy in such a way that the net used bandwidth tends towards the specified value. You can adjust this setting so that each map task is throttled to consume only the specified bandwidth so that the net used bandwidth tends towards the specified value.</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 4 connections.</p> <ul style="list-style-type: none"> If you set the value to 1 or more, Replication Manager uses multi-threading with the number of connections specified. If you set the value to 0 or fewer, Replication Manager uses single threading and a single connection.
Replication Option	Choose Metadata and Data to replicate metadata and data in files and directories. Otherwise, choose Metadata only to replicate the metadata of files and directories
Directory for metadata file	Enter / or a valid folder path in the target cluster to save the metadata file. If the field is empty or if the specified folder does not exist, Replication Manager creates a new folder.
Force Overwrite	<p>Select to overwrite data in the destination metastore if incompatible changes are detected.</p> <p>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. If you do not choose the option and the Hive replication process detects incompatible changes on the source cluster, Hive 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>
Invalidate Impala Metadata on Destination	<p>Choose the option to run the Impala <code>INVALIDATE METADATA</code> 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.</p> <p> Important: Do not select this option if the source and target clusters are Auto-TLS-enabled.</p> <p> Warning: 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.</p> <p> Note: If the source contains User Defined Functions (UDF), you must run the <code>INVALIDATE METADATA</code> statement manually and without any tables specified even if you configure the automatic invalidation.</p>
Replication Strategy	<p>Choose one of the following replication strategies:</p> <ul style="list-style-type: none"> Static distributes file replication tasks among the mappers up front to achieve an uniform distribution based on the file sizes. Dynamic distributes the file replication tasks in small sets to the mappers, and as each mapper completes its tasks, it dynamically acquires and processes the next set of unallocated tasks. <p>The default replication strategy is Dynamic.</p>
MapReduce Service	Choose the MapReduce or YARN service.
Log Path	Enter an alternate path for the logs, if required.

Option	Description
Error Handling	<p>Select the following options as necessary:</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> Note: 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 - 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 not selected by default. • Abort on Snapshot Diff Failures - If a snapshot diff fails during replication, the replication policy uses a complete copy to replicate data. If you select this option, the policy aborts the replication when it encounters an error instead.

Option	Description
Preserve	<p>Choose the required options 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.</p> <p>When Permission is selected, and both the source and destination clusters support ACLs, replication preserves ACLs. Otherwise, ACLs are not replicated. When Extended attributes is selected, 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.)</p> <p>If you select one or more of the Preserve options and you are replicating to S3 or ADLS, the values of all of these items are saved in metadata files on S3 or ADLS. When you replicate from S3 or ADLS to HDFS, you can select which of these options you want to preserve.</p> <p> Note: To preserve permissions to HDFS, you must be running as a superuser on the destination cluster. Use the Run As Username option to set the username.</p>
Delete Policy	<p>Choose the required options to determine whether the 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 option. • Delete to Trash - If the HDFS trash is enabled, files are moved to the trash folder. This is not supported when replicating to S3 or ADLS. • Delete Permanently - Uses the least amount of space; use with caution.
Alerts	<p>Choose to generate alerts for various state changes in the replication workflow. You can choose to generate an alert On Failure, On Start, On Success, or On Abort of the replication job.</p> <p>You can configure alerts to be delivered by email or sent as SNMP traps. If alerts are enabled for events, you can search for and view the alerts on the Events tab, even if you do not have email notification configured. For example, if you choose Command Result that contains the Failed filter on the Diagnostics Events page, the alerts related to the On Failure alert for all the replication policies for which you have set the alert appear. For more information, see Managing Alerts and Configuring Alert Delivery.</p>

14. Click Create.

If you selected Run Now on the **Schedule** page, the replication job starts data replication after you click Create.

Results

The replication task appears on the **Replication Policies** page.

Managing Hive replication policy




After you create a Hive replication policy in CDP Private Cloud Data Services Replication Manager, you can perform and monitor various tasks related to the replication policy. You can view the job progress and replication logs. You can edit the advanced options to optimize a job run. You can suspend a job and also activate a suspended job. You can edit the replication policy as necessary.

About this task

On the **Replication Policies** page, you can perform the following actions and tasks on a replication policy and its jobs.

Procedure

- When you click Actions for a Hive replication policy, the following actions appear:

Action	Description
Edit*	<p>Change the replication policy options as required for non-expired policies that are in active or suspended state. Based on the schedule you choose, the replication policy replicates data.</p> <p>You can edit the replication policies to better align with changing requirements. For example, you might want to change the frequency of a policy depending on the data size and importance of the data being replicated.</p> <p> Note: A replication policy is associated with a cluster or a cluster pair, therefore you cannot change the clusters in the policy.</p> <p>Optionally, expand a replication policy on the Replication Policies page to edit the replication policy options which include frequency (start time cannot be modified if the policy has already started), queue name, maximum bandwidth, and maximum map slots.</p> <p> Tip: To optimize the replication policy performance, you can configure the queue name, maximum bandwidth, and maximum map slots as necessary.</p>
Delete	Deletes the replication policy permanently.
Suspend	Suspends a running replication policy. Activate the replication policy, if required.
View Log	<p>Download, copy, or open the log to track the job and to troubleshoot any issues.</p> <p> Tip: On the Overview Issues & Updates panel, the Job Status column shows the current job status. If the job failed, click Failed to view the log details about the job.</p>
Collect diagnostic bundle	Generates a diagnostic bundle for the replication policy. You can download the bundle as a ZIP file to your machine.
<p>To view and use the replication policies with an empty name in Replication Manager, you must understand the following implementation:</p> <ul style="list-style-type: none"> If the Cloudera Manager API version is lower than 51, an existing replication policy with an empty name can be used and updated. However, if you edit the replication policy and provide a name for the replication policy in versions higher or equal to 51, you must ensure that the name conforms to the validation rules. If the Cloudera Manager API version is higher or equal to 51, it is mandatory that you provide a non-empty unique name to the replication policy to continue using it. This is because API version 51 and higher enforces the validation rules on all the replication policies. <p>To pass the replication policy name validation, you must ensure that the replication policy name is unique. The name can contain letters, numbers, and the <code>_</code> / <code>-</code> characters. You must also ensure that it does not contain the characters <code>%</code> <code>:</code> <code>;</code> <code>\</code> nor any character that is not ASCII printable, which includes the ASCII characters less than 32 and the ASCII characters that are greater than or equal to 127.</p>	

- When you expand the policy details, the **Job History** panel appears.

You can view the following details on the panel:

- Previous jobs, current job, and one future scheduled job if any.
- Job details which include:

Job details	Description
Started	Timestamp when the job started.
Ended	Timestamp when the job ended.
Duration	Time taken to complete the job.
Tables	Number of imported or exported tables.
Progress	Current status of a running job.

Job details	Description
Expected	Remaining number of files and bytes expected to be copied for a running job.
Copied	Number of files and bytes copied for a running job and completed job.
Failed	Number of files and bytes that failed to be copied for a completed job.
Deleted	Number of files deleted for a completed job.
Skipped	Remaining number of files and bytes skipped from copying for a running job and complete job.

c) Click Actions to:

- Abort the job.
- Re-run an aborted or failed job.
- View Log for the job. You can download, copy, or open it to track the job and to troubleshoot any issues for the job.
- When you click a job on the **Job History** panel, the following tabs appear:

Tab	Description
General	Shows the following job details: <ul style="list-style-type: none"> • Started at timestamp • Duration taken to complete the job • HDFS Replication Report to download the job statistics in CSV format • Hive Replication Report to download the job statistics in CSV format • Hive Export/Import is the number of external tables exported or imported using Hive replication. • Number of Errors encountered during the replication job. • Impala UDFs is the number of tables exported or imported using Impala. • Job status Message.
Command Details	Shows the details about the commands that ran on the source Cloudera Manager for the job, along with the timestamp.
Setup Error	Shows the stack trace for the commands that ran on the source Cloudera Manager for the failed job.

- You can download the following CSV reports from the **General HDFS Replication Report** field to track the replication jobs and to troubleshoot issues:

Report	Description
Listing	Lists all the files and directories copied during the replication job.
Status	Shows the complete status report of each file as: <ul style="list-style-type: none"> • an Error occurred and the file was not copied. • a Deleted file. • an up-to-date file for which the replication was Skipped.
Error Status	Status report of all the copied files with errors. Each file shows the status, path, and message for the copied files with errors.
Skipped Status	Status report of all skipped files. Each file lists the status, path, and message for the databases and tables that were skipped.
Deleted Status	Status report of all deleted files. Each file lists the status, path, and message for the databases and tables that were deleted.
Performance	Summary report about the performance of the running replication job which includes the last performance sample for each mapper that is working on the replication job.

Report	Description
Full Performance	Performance report of the job which includes the samples taken for all mappers during the replication job.



Note: The reports are generated based on the source Cloudera Manager response. If the Cloudera Manager response is interrupted or is not handled as expected, corresponding error messages appear in HTML format in the reports.

- You can download the following CSV reports from the **General Hive Replication Report** field to track the replication jobs and to troubleshoot issues:

Report	Description
Hive Result	List of replicated tables.
Hive Performance	Performance report for Hive replication.

Ozone replication policies

Apache Ozone is a scalable, distributed, and high performance object store optimized for big data workloads and can handle billions of objects of varying sizes. Ozone storage is co-located on HDFS. You can create Ozone replication policies to replicate data in Ozone buckets between clusters with CDP Private Cloud Base 7.1.8 or higher using Cloudera Manager 7.7.1 or higher.

Cloudera supports the following types of Ozone storage:

- Object store buckets (OBS), which are storage buckets where all the keys are written into a flat namespace and can be accessed using S3 interface provided by Ozone.
- File System Optimization (FSO), which are Hadoop-compatible file system buckets where the rename and delete operations on the directories are atomic. These buckets can be accessed using Filesystem APIs and S3 interfaces.
- Legacy buckets, which are Ozone buckets created prior to CDP Private Cloud Base 7.1.8 and use the Ozone File System (ofs) protocol or scheme.

You can use Ozone replication policies to replicate or migrate the required Ozone data to another cluster to run load-intensive workloads, back up data, or for backup-restore use cases.

Ozone replication policies support data replication between:

- FSO buckets in source and target clusters using ofs protocol.
- legacy buckets in source and target clusters using ofs protocol.



Note:

- If one or both of the source and destination buckets is a legacy bucket, then the `ozone.om.enable.filesystem.paths` flag (cluster-level configuration property) in the `ozone-site.xml` file must be enabled on the cluster(s) with the legacy bucket.
- Ozone replication uses `ofs` by default to replicate FSO or LEGACY buckets.
- OBS buckets in source and target clusters that support S3A filesystem using the S3A scheme or replication protocol.

How Ozone replication works

Ozone snapshots are enabled for all the buckets and volumes. If the incremental replication feature is enabled on the source and target clusters, to replicate Ozone data you can choose one of the following methods during the Ozone replication policy creation process:

Full file listing

By default, the Ozone replication policies use the full file listing method which takes a longer time to replicate data. In this method, the first Ozone replication policy job run is a bootstrap job; that

is, all the data in the chosen buckets are replicated. During subsequent replication policy runs, Replication Manager performs the following high-level steps:

1. Lists all the files.
2. Performs a checksum and metadata check on them to identify the relevant files to copy. This step depends on the advanced options you choose during the replication creation process. During this identification process, some unchanged files are skipped if they do not meet the criteria set by the chosen advanced options.
3. Copies the identified files from the source cluster to the target cluster.

Incremental only

In this method, the first replication policy job run is a bootstrap job, and subsequent job runs are incremental jobs.

To perform the incremental job, Replication Manager leverages Ozone snapshots and the snapshot-diff capability to generate a diff report. The diff report contains the changed or new data from the source cluster. The subsequent replication policy replicates data based on the diff report.

Incremental with fallback to full file listing

In this method, the first replication policy job run is a bootstrap job, and subsequent job runs are incremental jobs. However, if the snapshot-diff fails during a replication policy job run, the next job run is a full file listing run. After the full file listing run succeeds, the subsequent runs are incremental runs. This method takes a longer time to replicate data if the replication policy job falls back to the full file listing method.

Preparing clusters for Ozone replication policies

You must prepare the clusters, create buckets in the target cluster, and configure additional configurations for OBS bucket replication before you create Ozone replication policies.

About this task

Complete the following checklist to ensure that the clusters are ready to be used in an Ozone replication policy:

Procedure

- Have you added the source and target clusters on the [Management Console Clusters](#) page?

For more information, see [Adding clusters to a CDP Private Cloud Data Services deployment](#).



Important: Consider the following limitations before you add clusters:

- Only users with admin and poweruser roles can add clusters to use in Replication Manager.
- When Auto-TLS is enabled for a cluster, you must use the Cloudera Manager hostname to add a cluster.
- If a Cloudera Manager instance manages multiple clusters, you cannot include those clusters for replication in CDP Private Cloud Data Services Replication Manager.
- Have you created the bucket on the target cluster of the same type as the bucket on the source cluster from which the replication policy replicates data?



Tip: Create a volume and then the bucket. For more information, see [Managing storage elements using CLI](#).

The following sample commands create a volume and an FSO bucket:

```
ozone sh volume create o3://ozone1/vol1
ozone sh bucket create o3://ozone1/vol1/buck1 --layout FILE_SYSTEM_OPTIMIZED
```

- Are the additional configurations required for OBS bucket replication configured when the source bucket is an OBS bucket?

For more information, see [Configuring properties for OBS bucket replication](#).

- Do you need to replicate data securely? If so, ensure that the SSL/TLS certificate exchange between two Cloudera Manager instances that manage source and target clusters respectively is configured. For more information, see [Configuring SSL/TLS certificate exchange between two Cloudera Manager instances](#) on page 35.
- Is Kerberos enabled on both the clusters? If so, perform the following steps:
 - Configure a user with permissions to access HDFS and Ozone.
 - Run the following command to add the group name of the user (For example, the group name bdr) to the Ozone service configuration in target Cloudera Manager:

```
sudo usermod -a -G om bdr
```



Important: If Kerberos is enabled on both the clusters, you must run the `kinit -kt [***PATH**]/[***TO**]/ozone.keytab` command (the absolute path to the Ozone service's keytab) before you run any Ozone commands. For example, `kinit -kt /.../ozone.keytab om/[***PRINCIPAL**]/[***REALM.SAMPLE**]`.


- Is Ranger enabled on the source cluster? If so, you must:
 - complete the following steps on the Ranger UI from source Cloudera Manager:
 - Log into Ranger UI from source Cloudera Manager.
 - Click `cm_ozone` on the **Service Manager** page.
 - Add the user (that you configured in the previous step) to the all - volume, bucket, key, all - volume, and all - volume, bucket policy names, and then set the groups for this policy as public.
 - complete the following steps for the Ranger service in source Cloudera Manager:
 - Go to the source Cloudera Manager Clusters *Ranger service* Configuration tab.
 - Locate the Ranger KMS Server with KTS Advanced Configuration Snippet (Safety Valve) for `conf/kms-site.xml` property.
 - Add the following key-value pairs:
 - `hadoop.kms.proxyuser.om.hosts=*`
 - `hadoop.kms.proxyuser.om.groups=*`
 - `hadoop.kms.proxyuser.om.users=*`
 - Save the changes.
 - Restart the Ranger service for the changes to take effect.

Configuring properties for OBS buckets to use in Ozone replication policies

Before you replicate OBS buckets, you must configure additional properties that assist Ozone replication policies to replicate data in OBS buckets.


Procedure

- Add the key-value pairs in the following table to the Ozone Client Advanced Configuration Snippet (Safety Valve) property in the `ozone-site.xml` file in the source cluster:

Property	Action to take
<code>fs.s3a.endpoint</code>	Enter the same value as in Ozone S3 gateway web UI as the source cluster.  Tip: The source and target cluster have their own S3A endpoint URL.
<code>hadoop.tmp.dir</code>	Enter the temporary directory on the target cluster to buffer file uploads.
<code>fs.s3a.secret.key</code>	See Step 3 to get the required value.

Property	Action to take
fs.s3a.access.key	See Step 3 to get the required value.
ozone.om.snapshot.load.native.lib	Enter false. Incremental Ozone replication policy runs use snapshot-diff operation. This property ensures that the replication policy run is not affected if the snapshot-diff operation goes down during the replication policy run.

2. Add the key-value pairs in the following table to the Ozone Client Advanced Configuration Snippet (Safety Valve) property in the ozone-site.xml file in the target cluster:

Property	Action to take
fs.s3a.endpoint	Enter the same value as in Ozone S3 gateway web UI as the target cluster.  Tip: The source and target cluster have their own S3A endpoint URL.
fs.s3a.secret.key	See Step 3 to get the required value.
fs.s3a.access.key	See Step 3 to get the required value.
fs.s3a.change.detection.version.required	Set to false.
fs.s3a.change.detection.mode	Enter none.
fs.s3a.path.style.access	Enter true.
ozone.om.snapshot.load.native.lib	Enter false. Incremental Ozone replication policy runs use snapshot-diff operation. This property ensures that the replication policy run is not affected if the snapshot-diff operation goes down during the replication policy run.

3. If Kerberos is enabled on the source and target cluster, run the `ozone s3 getsecret --om-service-id=serviceId` command to get the secret and access key. Otherwise, enter any arbitrary value for the secret and access key.

You can store the keys in a credstore such as JCEKS for non Auto-TLS clusters. After you store the keys, perform the following steps:

- Configure the credstore file path for the `hadoop.security.credential.provider.path` property in the `ozone-site.xml` file. For more information, see [Using DistCp with Amazon S3](#).
- Add the `HADOOP_CREDSTORE_PASSWORD` parameter to the YARN Service Environment Advanced Configuration Snippet (Safety Valve) property for the YARN service in source Cloudera Manager.



Note: If no password is set, enter none for the property.

4. The `/s3v` volumes store S3 buckets. By default, you can access the buckets in `/s3v` volumes using S3 interface. To access other buckets through the S3 interface, you must create a “symbolic linked” bucket. You can use the ‘symbolic linked’ bucket in Ozone replication policies.

Configure the required OBS buckets as S3-compatible buckets using the following commands before you use it in Ozone replication policies:

- `ozone sh volume create /s3v`
- `ozone sh volume create /[***vol_name***]`
- `ozone sh bucket create /[***vol_name***]/[***bucket_name***]`
- `ozone sh bucket link /[***vol_name***]/[***bucket_name***] /s3v/[***symbolic_linked_bucket_name***]`

5. Import the S3G CA certificate from the cluster to the local JDK path using the following commands:

- a) Run the `keytool -importkeystore -destkeystore [***jdk_cacerts_location***] -srckeystore [***cm-auto-global_truststore.jks location***] -srcalias [***cm_alias_on_src_cm***]` command on all the hosts of the *source* Cloudera Manager.

For example, `keytool -importkeystore -destkeystore /usr/java/default/lib/security/cacerts -srckeystore /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_truststore.jks -srcalias cmrootca-0`

- b) Run the following commands on all the hosts of the target Cloudera Manager:

1. `keytool -importkeystore -destkeystore [***jdk_cacerts_location***] -srckeystore [***cm-auto-global_truststore.jks location***] -srcalias [***cm_alias_on_src_cm***]`
2. `keytool -importkeystore -destkeystore [***jdk_cacerts_location***] -srckeystore [***cm-auto-global_truststore.jks location***] -srcalias [***cm_alias_on_dest_cm***]`

For example,

```
keytool -importkeystore -destkeystore /usr/java/default/lib/security/cacerts
-srckeystore /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_tru
ststore.jks -srcalias cmrootca-0
keytool -importkeystore -destkeystore /usr/java/default/lib/security/ca
certs
-srckeystore /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_tr
uststore.jks -srcalias cmrootca-1
```

Creating an Ozone replication policy

You can create an Ozone replication policy in CDP Private Cloud Data Services Replication Manager to replicate Ozone data between CDP Private Cloud Base 7.1.8 or higher clusters.

Procedure


1. Log into CDP Private Cloud Data Services, and click Replication Manager.
2. Click Replication Policies.
3. On the **Replication Policies** page, click Create Policy.
4. On the **General** page, choose or enter the following information:

Option	Action to take
Ozone	Choose to create an Ozone replication policy.
Policy Name	Enter a unique name for the replication policy.
Description	Optional. Enter a brief description about the replication policy.

5. Click Next.

6. On the **Select Source** page, choose or enter the following information:

Option	Action to take
Source cluster	Choose the source cluster.
Run As Username (on source)	Optional. The replication policy uses the <i>Default</i> username to replicate data. If you are using a kerberized cluster, enter the required username. The replication policy uses this username to replicate the data in the kerberized cluster.

Option	Action to take
Source path	<p>Choose one of the following path types depending on the Ozone storage:</p> <ul style="list-style-type: none"> FSO (FileSystemOptimized) to FSO - Enter the volume and bucket names in the source cluster. OBS (ObjectStore) to OBS - Enter the bucket name in the source cluster. <p> Important: Complete the steps in Configuring properties for OBS buckets to use in Ozone replication policies on page 24 before you use this option.</p> <ul style="list-style-type: none"> Full Path - Enter the path to the bucket in the ofs://[***ozone_service_id***]/[***volume_name***]/[***bucket_name***] or s3a://[***bucket_name***] format to replicate data between FSO or OBS buckets respectively. A bucket subpath can also be specified.

7. Click Next.

8. On the **Select Destination** page, choose or enter the following information:




Option	Action to take
Destination cluster	Choose the target cluster.
Destination Volume Destination Bucket	Options appear depending on the path type you choose on the Select Source page. Enter the volume and bucket in the target cluster as necessary.
Run as Username	<p>Optional. The replication policy uses the <i>Default</i> username to replicate the data.</p> <p>Enter another username if you want the replication policy to use it to replicate data.</p>


9. Click Validate Policy.

Replication Manager verifies whether the details provided are correct.

10. Click Next.

11. On the **Schedule** page, choose or enter the following information:


Option	Action to take
Run Now	Choose to initiate data replication after the replication policy creation is complete. Choose the frequency to replicate data periodically, if required.
Schedule Run	<p>Choose to run the replication policy to replicate data at a later time. Choose the date and time for the first run, and then choose the frequency to replicate data periodically.</p> <p> Tip: On the Replication Policies page, click   to change the timezone.</p>

Option	Action to take
Frequency	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> Does Not Repeat. Custom - In the Custom Recurrence dialog box, choose the time, date, and the frequency to run the policy. <p> Note: Ensure that the frequency in a schedule enables a job to finish before the next job starts. Also, ensure that the jobs based on the same policy do not overlap. If a job is not completed before another job starts, the second job does not run and the job status appears as Skipped. If a job is consistently skipped, you might need to modify the frequency of the job.</p>

12. Click Next.

13. On the **Additional Settings** page, enter the attribute values, as necessary:

Option	Description
YARN Queue Name	Enter the name of the YARN queue for the cluster to which the replication job is submitted if you are using Capacity Scheduler queues to limit resource consumption. The default value for this field is <i>default</i> .
Maximum Maps Slots	Set the maximum number of map tasks (simultaneous copies) per replication job. The default value is 20.
Maximum Bandwidth	<p>Adjust this setting so that each map task is throttled to consume only the specified bandwidth. The default value for the bandwidth is 100 MB per second for each mapper.</p> <p>Each map task (simultaneous copy) is restricted to consume only the specified bandwidth. This is not always exact. The map throttles back its bandwidth consumption during a copy in such a way that the net used bandwidth tends towards the specified value. You can adjust this setting so that each map task is throttled to consume only the specified bandwidth so that the net used bandwidth tends towards the specified value.</p>
Listing type	<p>Choose one of the following replication methods to replicate Ozone data:</p> <ul style="list-style-type: none"> Full file listing. Incremental only Incremental with fallback to full file listing <p>To understand how each method works, see Ozone replication policies.</p> <p>This option appears only if the incremental replication feature is enabled on the source and target clusters.</p>
Replication Strategy	<p>Choose one of the following replication strategies:</p> <ul style="list-style-type: none"> Static distributes file replication tasks among the mappers up front to achieve an uniform distribution based on the file sizes. Dynamic distributes the file replication tasks in small sets to the mappers, and as each mapper completes its tasks, it dynamically acquires and processes the next set of unallocated tasks. <p>The default replication strategy is Dynamic.</p>
MapReduce Service	Choose the MapReduce or YARN service.
Log Path	Enter an alternate path for the logs, if required.

Option	Description
Error Handling	<p>Select the following options as necessary:</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> Note: 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. • 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 not selected by default.
Delete Policy	<p>Choose the required options to determine whether the 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 option. • Delete to Trash - If the HDFS trash is enabled, files are moved to the trash folder. This is not supported when replicating to S3 or ADLS. • Delete Permanently - Uses the least amount of space; use with caution.

Option	Description
Alerts	<p>Choose to generate alerts for various state changes in the replication workflow. You can choose to generate an alert On Failure, On Start, On Success, or On Abort of the replication job.</p> <p>You can configure alerts to be delivered by email or sent as SNMP traps. If alerts are enabled for events, you can search for and view the alerts on the Events tab, even if you do not have email notification configured. For example, if you choose Command Result that contains the Failed filter on the Diagnostics Events page, the alerts related to the On Failure alert for all the replication policies for which you have set the alert appear. For more information, see Managing Alerts and Configuring Alert Delivery.</p>

14. Click Create.

If you selected Run Now on the **Schedule** page, the replication job starts data replication after you click Create.

Results

The replication task appears on the **Replication Policies** page. If you selected Run Now on the **Schedule** page, the replication job starts data replication when you click Create.

Managing Ozone replication policy



After you create an Ozone replication policy in CDP Private Cloud Data Services Replication Manager, you can perform and monitor various tasks related to the replication policy. You can view the job progress and replication logs. You can suspend a job and also activate a suspended job.


About this task

On the **Replication Policies** page, you can perform the following actions and tasks on a replication policy and its jobs:

Procedure

- Click Actions for the following options:

Action	Description
Edit*	<p>Change the replication policy options as required for non-expired policies that are in active or suspended state. Based on the schedule you choose, the replication policy replicates data.</p> <p>You can edit the replication policies to better align with changing requirements. For example, you might want to change the frequency of a policy depending on the data size and importance of the data being replicated.</p> <p> Note: A replication policy is associated with a cluster or a cluster pair, therefore you cannot change the clusters in the policy.</p> <p>Optionally, expand a replication policy on the Replication Policies page to edit the replication policy options which include frequency (start time cannot be modified if the policy has already started), queue name, maximum bandwidth, and maximum map slots.</p> <p> Tip: To optimize the replication policy performance, you can configure the queue name, maximum bandwidth, and maximum map slots as necessary.</p>
Rerun	Runs the replication policy.
Delete	Deletes the replication policy permanently.

Action	Description
Suspend	Suspends a running replication policy. Activate a suspended replication policy, as necessary.
View command details	Opens the latest replication policy job page. The steps and substeps appear in a tree view. The failed steps are expanded by default, showing the last 15 lines of the log. You can also use this option on the Overview Issues & Updates panel.  Tip: You can view the complete log for all the jobs on the target cluster Cloudera Manager Running Commands page.
Collect diagnostic bundle	Generates a diagnostic bundle for the replication policy. You can download the bundle as a ZIP file to your machine.
<p>To view and use the replication policies with an empty name in Replication Manager, you must understand the following implementation:</p> <ul style="list-style-type: none"> If the Cloudera Manager API version is lower than 51, an existing replication policy with an empty name can be used and updated. However, if you edit the replication policy and provide a name for the replication policy in versions higher or equal to 51, you must ensure that the name conforms to the validation rules. If the Cloudera Manager API version is higher or equal to 51, it is mandatory that you provide a non-empty unique name to the replication policy to continue using it. This is because API version 51 and higher enforces the validation rules on all the replication policies. <p>To pass the replication policy name validation, you must ensure that the replication policy name is unique. The name can contain letters, numbers, and the _ / - characters. You must also ensure that it does not contain the characters % . ; \ nor any character that is not ASCII printable, which includes the ASCII characters less than 32 and the ASCII characters that are greater than or equal to 127.</p>	

- Expand the policy details to view the **Job History** panel.

The following details appear on the panel:

- Previous jobs, current job, and one future scheduled job if any.
- Job details which include:

Job details	Description
Started	Timestamp when the job started.
Ended	Timestamp when the job ended.
Duration	Time taken to complete the job.
Progress	Current status of a running job.

- Click a job on the **Job History** panel to view the following tabs:

Tab	Description
General	Shows the following job details: <ul style="list-style-type: none"> Started at timestamp Duration to complete the job HDFS Replication Report to download the job statistics in CSV format Job status Message
Command Details	Shows the job steps along with the timestamp.

- Download the following CSV reports from the General HDFS Replication Report field to track the replication jobs and to troubleshoot issues:

Report	Description
Listing	Lists all the files and directories copied during the replication job.
Status	Status report of the files. Shows the following status for each file: <ul style="list-style-type: none"> an Error occurred and the file was not copied. a Deleted file. an up-to-date file for which the replication was Skipped.

Report	Description
Error Status	Status report of all the copied files with errors. Each file shows the status, path, and message for the copied files with errors.
Skipped Status	Status report of all skipped files. Each file lists the status, path, and message for the databases and tables that were skipped.
Deleted Status	Status report of all deleted files. Each file lists the status, path, and message for the databases and tables that were deleted.
Performance	Summary report about the performance of the running replication job which includes the last performance sample for each mapper that is working on the replication job.
Full Performance	Performance report of the job which includes the samples taken for all mappers during the replication job.



Note: The reports are generated based on the source Cloudera Manager response. If the Cloudera Manager response is interrupted or is not handled as expected, the *Replication report not found for command with id: {command_id}* error appears.

Appendix

You must configure the kerberized clusters with the necessary configurations in Cloudera Manager before you create replication policies in CDP Private Cloud Data Services Replication Manager to replicate data between CDP Private Cloud Base 7.1.8 or higher clusters. You must also add the required user and group to all the hosts in the source and destination clusters in Cloudera Manager before you replicate data from unsecure to secure clusters using replication policies.

Replication between clusters using Kerberos authentication

After you enable Kerberos authentication on the source and destination clusters, you must perform some additional steps to enable replication on the clusters. The additional steps include opening the required ports on the clusters, checking the realm names to avoid conflicts when running replication jobs, and configuring the source and target clusters.

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

- Secure source to a secure destination.
- Insecure source to an insecure destination.
- Insecure source to a secure destination when the following requirements are met:
 - 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 configuration steps in *Configuring user to replicate from unsecure to secure clusters* must be complete.

To enable Kerberos authentication, see [Enabling Kerberos authentication for CDP](#). To enable replication between clusters using Kerberos authentication, see *Configuring kerberized clusters for replication*.



Note:

- 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 configuration, and are not related to Kerberos realm trust relationships.
- If you are using a standalone DistCp job between clusters in different Kerberos realms, you must configure a realm trust.

Configuring kerberized clusters for replication

To replicate data between CDP Private Cloud Base 7.1.8 or higher clusters using Kerberos authentication, you must ensure the required ports are open and the clusters are configured as required.

Procedure

1. Open the port used for the Kerberos KDC Server and KRB5 services on all the hosts on the destination cluster. By default, this is port 88.
2. Use one of the following configurations to ensure that the realm names do not create conflicts while 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.
 - 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 might result in principals that have equal permissions in both environments. Make sure that permissions are set appropriately for each type of environment.
 - Replication fails 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.
3. Configure the following steps on the source and destination clusters so that Replication Manager can replicate data across clusters in two different realms to set up trust between those clusters:
 - a. On the hosts in the *destination* cluster, ensure that the `krb5.conf` file (typically located at the `/etc/krb5.conf` location) on each host has the following information:
 1. 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
}
```

2. Realm mapping for the *source* cluster domain 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
```



Note: If you have a scenario where the hostname(s) are inconsistent, ensure that all those hosts are covered in a similar manner as seen in the `domain_realm` section in the Cloudera Manager Host All Hosts section.

b. On the *destination* cluster, perform the following steps:

1. Search for the Trusted Kerberos Realms property on the *HDFS service* Configuration tab.
2. Enter the source cluster realm name, and click Save Changes.
3. Search for the Domain Name(s) field in the Administration Settings section.
4. Enter the domain or host names you want to map to the destination cluster KDC. Add as many entries as you need.

The entries in this property are used to generate the domain_realm section in the krb5.conf file.

5. Click Save Changes.
6. Remove domain_realm entries in the Advanced Configuration Snippet (Safety Valve) for remaining krb5.conf property.

Configuring user to replicate from unsecure to secure cluster

Configure a user (on all the hosts on both the source and destination clusters) that Replication Manager can use to replicate data from an unsecure cluster (one that does not use Kerberos authentication) to a secure cluster (a cluster that uses Kerberos). If required, specify this user in the "Run As Username" field when you create a replication policy. Replication Manager does not support replicating from a secure cluster to an unsecure cluster, or a mixture of secure and unsecure source clusters in replication scenarios where a destination cluster has multiple source clusters.

Procedure

1. Add a user on a host in the source or destination cluster with the `sudo -u hdfs hdfs dfs -mkdir -p /user/r/[***username***]` command.

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 `sudo -u hdfs hdfs dfs -chown [***username***] /user/username` command.

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 `groupadd supergroup` command.
4. Add the user you created in step 1 to the group you created with the `usermod -G supergroup [***username***]` command.

For example, the following command adds milton to the group named supergroup:

```
usermod -G supergroup milton
```

5. Repeat step 1 through step 4 for all the 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 these steps, ensure that you specify this user in the Run As Username field when you create a replication policy to replicate data from an unsecure source cluster to a secure target cluster.

Kerberos connectivity test

As part of the 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.

Kerberos connectivity test 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 the clusters are 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.

Configuring SSL/TLS certificate exchange between two Cloudera Manager instances

The Replication Manager configures replication peers between two clusters before running the replication job. You can manually set up an SSL/TLS certificate exchange between two Cloudera Manager instances that manage source and target cluster respectively. Replication Manager uses this information to set up the peers for secure data replication.

About this task

Replication Manager supports Cloudera Manager high availability functionality only after you manually configure the SSL/TLS certificate exchange.

When the source Cloudera Manager is configured for high availability and is Auto-TLS enabled, the certificate exchange is initiated from the source cluster to the target cluster where the certificate is exported from the load balancer node of the source cluster.



Important: The following sample commands use the *open-jdk-11* Java version. Use the Java version that you use in CDP clusters in these commands.

Procedure

1. Go to the truststore location in *source* Cloudera Manager, and perform the following steps:

- a) List the contents of the keystore file and password using the `[***keytool path***] -list -keystore [***truststore JKS file location ***] -storepass [***truststore password***]` command.

For example, `/usr/lib/jvm/java-openjdk-11/bin/keytool -list -keystore /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_truststore.jks -storepass [***truststore password***]`



Tip:

- The keytool path can be located in various locations including the keytool itself. For example, it can be located in `/usr/lib/jvm/java-openjdk-11/bin/keytool` or `/usr/java/default/bin/keytool`.
- You can locate the truststore password using the `cat /etc/hadoop/conf/ssl-client.xml` command. You can enter the SSL password for the `/etc/hadoop/conf/ssl-client.xml` file when prompted.
- Alternatively, you can also run the following commands instead of the command in Step a:

```
export JAVA_HOME=[***keytool location***]

export TRUSTSTORE_JKS=[***truststore JKS file location***]

export TRUSTSTORE_PASSWORD=[***password in the ssl-client.xml
file***]
$JAVA_HOME/keytool -list -keystore
$TRUSTSTORE_JKS -storepass
$TRUSTSTORE_PASSWORD
```

- b) Export the certificate contents in the host to a file using the `[***keytool***] -exportcert -keystore [***truststore JKS file location ***] -alias [***cm_alias_on_src_cm***] -file ./[***TXT file, for example: source-cert.txt***] -storepass [***truststore_password***]` command.

For example,

```
/usr/java/default/bin/keytool -exportcert -keystore /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_truststore.jks -alias cmrootca-0 -file
./source-cert.txt -storepass [***truststore_password***]
```


- c) Copy the text file to all the hosts of the *target* cluster Cloudera Manager securely using the `scp -i [***PEM file***] [***TXT file - source-cert.txt***] root@[***host_ip***]:/home/` command.
- d) Import the certificate into the keystore file on all the hosts of the *target* cluster Cloudera Manager using the `[***keytool***] -importcert -noprompt -v -trustcacerts -keystore [***truststore JKS file location ***] -alias [***cm_alias_on_dest_cm***] -file ./[***TXT file - source-cert.txt***] --storepass [***truststore_password***]` command.

For example, `/usr/java/default/bin/keytool -importcert -noprompt -v -trustcacerts -keystore /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_truststore.jks -alias cmrootca-1 -file ./source-cert.txt --storepass [***truststore_password***]`

2. Go to the truststore location in *target* Cloudera Manager, and perform the following steps:

- a) List the contents of the keystore file and password using the `[***keytool path***] -list -keystore [***truststore JKS file location ***] -storepass [***truststore password***]` command.
- b) Export the certificate contents in the host to a file using the `[***keytool***] -exportcert -keystore [***truststore JKS file location ***] -alias [***cm_alias_on_dest_cm***] -file ./[***TXT file, for example: dest-cert.txt***] -storepass [***truststore_password***]` command.
- c) Copy the text file to all the hosts of the *source* cluster Cloudera Manager securely using the `scp -i [***PEM file***] [***TXT file - dest-cert.txt***] root@[***host_ip***]:/home/` command.
- d) Import the certificate into the keystore file on all the hosts of the *source* Cloudera Manager using the `[***keytool***] -importcert -noprompt -v -trustcacerts -keystore [***truststore JKS file location ***] -alias`

`[***cm_alias_on_src_cm***] -file ./[***TXT file - dest-cert.txt***] --storepass [***truststore_password***]`
command.

3.  **Note:** Perform this step only for Ozone replication policies.

Import the S3G CA certificate from the cluster to the local JDK path using the following commands:

- a) Run the `keytool -importkeystore -destkeystore [***jdk_cacerts_location***] -srckeystore [***cm-auto-global_truststore.jks location***] -srcalias [***cm_alias_on_src_cm***]` command on all the hosts of the source Cloudera Manager.

For example, `keytool -importkeystore -destkeystore /usr/java/default/lib/security/cacerts -srckeystore /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_truststore.jks -srcalias cmrootca-0`

- b) Run the following commands on all the hosts of the target Cloudera Manager:

1. `keytool -importkeystore -destkeystore [***jdk_cacerts_location***] -srckeystore [***cm-auto-global_truststore.jks location***] -srcalias [***cm_alias_on_src_cm***]`
2. `keytool -importkeystore -destkeystore [***jdk_cacerts_location***] -srckeystore [***cm-auto-global_truststore.jks location***] -srcalias [***cm_alias_on_dest_cm***]`



Tip: Enter the `security/jssecacerts` path for the `-destkeystore` attribute if the file exists. Otherwise, enter the `security/cacerts` path.

For example,

```
keytool -importkeystore -destkeystore /usr/java/default/lib/security/cac
erts
-srckeystore /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_tru
store.jks -srcalias cmrootca-0
keytool -importkeystore -destkeystore /usr/java/default/lib/security/ca
certs
-srckeystore /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_tr
uststore.jks -srcalias cmrootca-1
```



Note: If you do not complete Step 3 before you create and run an Ozone replication policy, an SSL certificate exception might appear during the file listing phase of the Ozone replication policy job run.