# DLM Administration

**Date of Publish:** 2019-02-28


**https://docs.hortonworks.com/**

# Contents

# Introduction

This *System Administrator Guide* provides information on how to setup, configure, and manage Data Lifecycle Manager (DLM) jobs to replicate data for small and large scale enterprise organisations.

## Purpose and scope

This guide is intended to assist System Administrators to manage the complete data replication life cycle, using on-premise and cloud environments.

## Audience and assumptions

The intended audience is composed of DLM System Administrators and community-based end-users. This document assumes that the reader has some experience installing and administering DataPlatform applications.

# Replication concepts

Data replication concepts involving DLM.

## Data Lifecycle Manager terminology

Data Lifecycle Manager (DLM) is a UI service that is enabled through DP Platform. From the DLM UI you can create and manage replication and disaster recovery policies and jobs.

| Term | Description |
|---|---|
| DLM App or Service | The web UI that runs on the DP platform host. The corresponding agent needs to be installed on the clusters. |
| DLM Engine | The agent required for DLM. Also referred to as the DLM Engine, this replication engine must be installed as a management pack on each cluster to be used in data replication jobs. The engine maintains, in a configured database, information about clusters and policies that are involved in replication. |
| Data center | The facility that contains the computer, server, and storage systems and associated infrastructure, such as routers, switches, and so forth. Corporate data is stored, managed, and distributed from the data center. In an on-premise environment, a data center is often composed of a single HDP cluster. However, a single data center can contain multiple HDP clusters. |
| IaaS cluster | A full HDP cluster on cloud VMs with Apache services running, such as HDFS, YARN, Ambari, Hiveserver2, Ranger, Atlas, and DLM Engine. Replication behavior is similar to on-premise cluster replication. The data is on local HDFS. |
| Cloud data lake or data lake | An HDP cluster on the cloud, using VMs, with data retained on cloud storage. A cloud data lake requires minimal services for metadata and governance, such as Hive metastore, Ranger, Atlas, and DLM Engine. |
| Cloud storage | Any storage retained in a cloud account, such as Amazon S3 web service. |

| Term | Description |
|---|---|
| On-premise cluster | A full HDP cluster in a data center, with Apache services running, such as HDFS, Yarn, HMS, hiveserver2, Ranger, Atlas and DLM Engine. Replication behavior is similar to IaaS cluster replication. The data is on local HDFS. |
| Policy | A set of rules applied to a replication relationship. The rules include which clusters serve as source and destination, the type of data to replicate, the schedule for replicating data, and so on. |
| Job | An instance of a policy that is running or has run. |
| Source cluster | The cluster that contains the source data that will be replicated to a destination cluster. Source data could be an HDFS dataset or a Hive database. |
| Destination cluster | The cluster to which an HDFS dataset or Hive database is replicated. |
| Target | The path on the destination cluster to which the HDFS dataset or Hive database is replicated. |

**Related Concepts**

DP Platform terminology

## Communicating within services in HDP

DLM App interacts with multiple services within HDP to process replication requests.

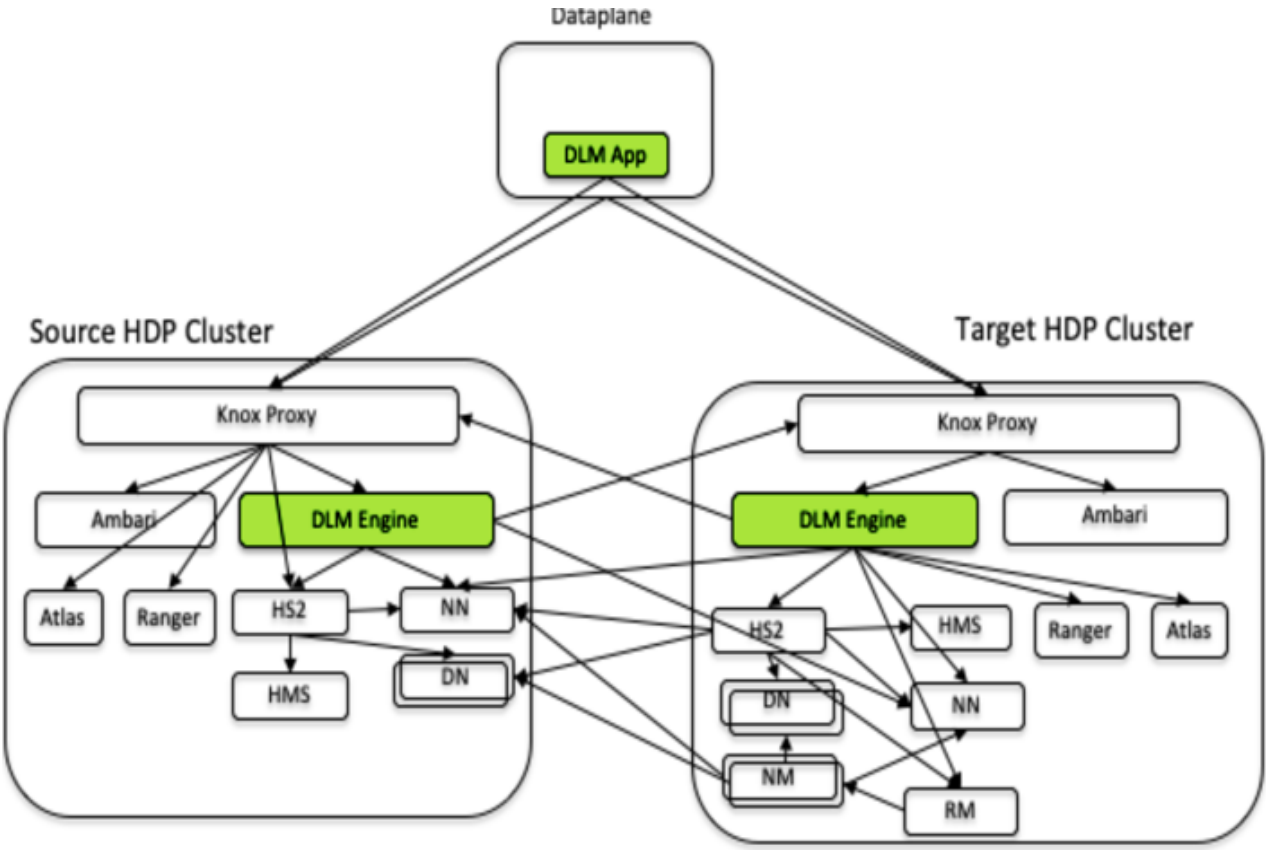


- **HS2**: Hiveserver2 (Hive)
- **HMS**: Hive Metastore Service (Hive)
- **NN**: Namenode (HDFS)
- **DN**: Datanode (HDFS)
- **RM**: Resource Manager (YARN)
- **NM**: Node Manager (YARN)
- **DLM App**: Runs in Dataplane host as a docker container. DLM App hosts DLM UI and DLM App APIs.
- **DLM Engine**: Agent service that runs in each HDP cluster, exposes REST APIs for DLM App.

## How Policies Work in Data Lifecycle Manager

In Data Lifecycle Manager, you create policies to establish the rules you want applied to your replication and disaster recovery jobs. The policy rules you set can include which cluster is the source and which is the destination, what data is replicated, what day and time the replication job occurs, the frequency of job execution, and bandwidth restrictions.

When scheduling how often you want a replication job to run, you should consider the recovery point objective (RPO) of the data being replicated; that is, what is the acceptable lag time between the active site and the replicated data on the destination. Data Lifecycle Manager supports a one-hour RPO: data is preserved up to one hour prior to the point of data recovery. To meet a one-hour RPO, you must consider how long it takes to replicate the selected data, how often the data is replicated, and network bandwidth capabilities.

As an example, if you have a set of data that you expect to take 15 minutes to replicate, then to meet a one-hour RPO, you would schedule the replication job to occur no more often than every 45 minutes, depending on network bandwidth.

- The first time you execute a job (an instance of a policy) with data that has not been previously replicated, Data Lifecycle Manager creates a new folder or database and bootstraps the data.

  During a bootstrap operation, all data is replicated from the source cluster to the destination. As a result, the initial execution of a job can take a significant amount of time, depending on how much data is being replicated, network bandwidth, and so forth. So you should plan the bootstrap accordingly.

  After initial bootstrap, data replication is performed incrementally, so only updated data is transferred. Data is in a consistent state only after incremental replication has captured any new changes that occurred during bootstrap.
- Achieving a one-hour Recovery Point Objective (RPO) depends on how you set up your replication jobs and the configuration of your environment:
  - Select data in sizes that replicate within 30 minutes.
  - Set replication frequency to 45 minutes or less.
  - Ensure that network bandwidth is sufficient, so that data can move fast enough to meet your RPO.
  - Consider the rate of change of data being replicated.

## Replication policy

You should take into consideration the following items when creating or modifying a replication policy.

**Data security**

- If using TDE for encryption, the entire source directory must be either encrypted or not encrypted, otherwise policy creation fails.
- If using an S3 cluster for your policy, your credentials must have been registered on the Cloud Credentials page.
- On destination clusters, the DLM Engine must have been granted write permissions on folders being replicated.
- Any user with access to the DLM UI has the ability to browse, within the DLM UI, the folder structure of any clusters enabled for DLM.

  Therefore, the DataPlane Admins and the Infrastructure Admins can view folders, files, and databases in the DLM UI that they might not have access to in HDFS. The **DataPlane Admin**, **Infrastructure Admin**, and **DLM Admin** cannot view from the DLM UI the content of files on the source or destination clusters. Nor do these administrators have the ability to modify or delete folders or files that are viewable from the DLM UI.

**Policy properties and settings**

- Ensure that the frequency is set so that a job finishes before the next job starts. Jobs based on the same policy cannot overlap. If a job is not completed before another job starts, the second job does not execute and is given the status Skipped. If a job is consistently skipped, you might need to modify the frequency of the job.
- Specify bandwidth per map, in MBps. Each map 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 bandwidth used tends towards the specified value.

**Cluster requirements**

- The target folder or database on the destination cluster must either be empty or not exist prior to starting a new policy instance.
- The clusters you want to include in the replication policy must have been paired already.
- With a single cluster, you can replicate data on-premise to cloud and vice-versa.

- With a single cluster, you cannot replicate data on-premise to on-premise and vice-versa.
- On the **Create Policy** page, the only requirement for clusters to display in the **Source Cluster** or **Destination Cluster** fields is that they are DLM-enabled. You must ensure that the clusters you select are healthy before you start a policy instance (job).

**Hive restrictions**

- ACID tables and storage handler-based tables (such as HBase) are currently not replicated.
- When creating a schedule for a Hive replication policy, you should set the frequency so that changes are replicated often enough to avoid overly large copies.

# UI overview

The **UI overview** describes various components of DLM app.

## Cluster Health panel

You can use the **Cluster Health** panel of the Overview page to view the total number of clusters enabled for Data Lifecycle Manager, the number that are healthy, the number for which a warning is issued, and the number that are unhealthy.

You can investigate the issues associated with clusters that have a warning or unhealthy status by navigating to the Ambari web UI.

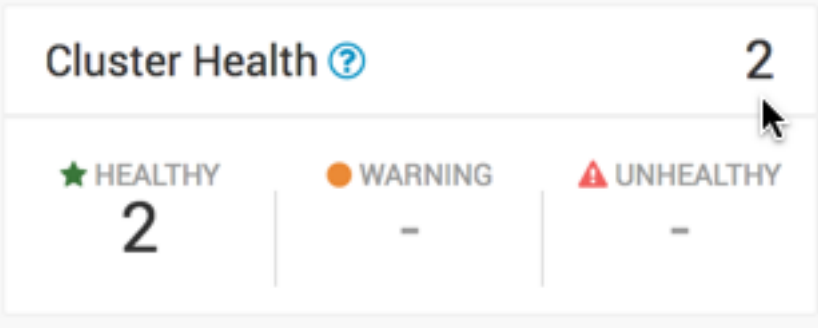


| | |
|---|---|
| **Healthy** | Specifies the total number of clusters currently available to run replication jobs. The DLM Engine can be reached and all services are running. |
| **Warning** | Specifies the total number of clusters for which remaining disk capacity is less than 10%.<br><br>If this value is greater than zero, you can click the number to open a table that displays the cluster name and remaining capacity. |
| **Unhealthy** | Specifies the total number of clusters for which at least one Apache Ambari service required for DLM (DLM Engine, HDFS, Apache Hive, or Apache Knox) is not started. If this value is greater than zero, you can click the number to open a table that displays the cluster name and the names of any Ambari services that have stopped. |

## Policies panel

You can use the **Policies panel** of the **Overview** page to view the total number of policies in use and their status.

| | |
|---|---|
| **Active** | Specifies policies with status of Submitted or Running. This item is not actionable |
| **Suspended** | Specifies policies that have been suspended by an administrator. This item is not actionable. |
| **Unhealthy** | Specifies policies associated with any cluster designated as Unhealthy in the **Cluster Health** panel. If the value is greater than zero, the number becomes clickable. You can click the number to display a table that contains the policy name, the names of the source and destination clusters, and which services are stopped on the source or destination cluster. |

## Jobs panel

You can use the **Jobs** panel of the **Overview** page to view the total number of running and failed jobs and their status.

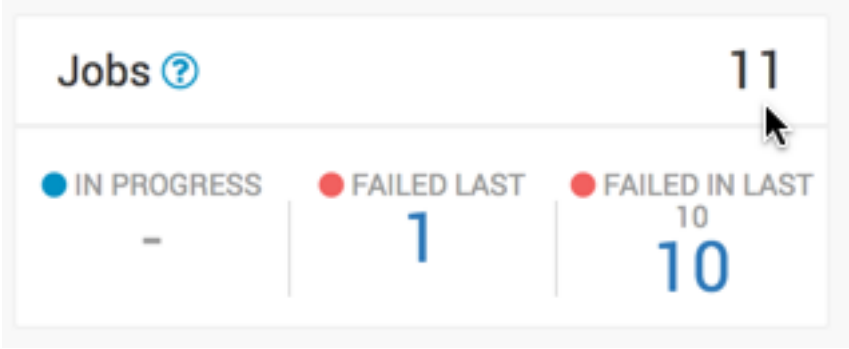


| | |
|---|---|
| **In Progress** | Specifies the number of jobs with the status Running. If the value is greater than zero, the number becomes clickable. You can click the number to apply a filter to the Issues and Updates table, so that the table displays only in-progress jobs. The filter label Jobs: In Progress appears above the table. Running jobs display as a blue dot in the Policy History column of the table. |
| **Failed Last** | Specifies the last job that completed with status Failed. If the value is greater than zero, the number becomes clickable. You can click the number to apply a filter to the Issues and Updates table, so that the table displays only policies for which the last job had a status of Failed. The filter label Jobs: Failed Last appears above the table. Failed jobs display as a red dot in the Policy History column of the table. |
| **Failed in Last 10** | Indicates the number of policies for which at least one of the last 10 jobs completed with status Failed. If the value is greater than zero, the number becomes clickable. You can click the number to apply a filter to the Issues and Updates table, so that the table displays only policies for which at least one job failed out of the last 10 jobs. The filter label Jobs: Failed Last appears above the table. |

| | |
|---|---|
| | Failed jobs display as a red dot in the Policy History column of the table. |

## Recent Issues panel

This panel shows the last four events with severity of warning, critical, or error. For each event, the panel shows the severity, the type, a message that includes the policy name and file icon, and the age of the event.

| | |
|---|---|
| **Severity icon** | Displays in orange for warning and red for critical or error. |
| **Event type** | Displays in bold text above the event message. The type can be succeeded, deleted, or suspended. |
| **Event message** | Displays as text under the event name. When an event is associated with a policy or policy instance (job), then the message text contains two items:<br>• **Policy name**: You can click this term to navigate to the Policies page with a preset filter that displays information about only the selected policy.<br>• **File icon**: You can move the cursor over the icon to display the text “View Log” and click the text to display log content for the associated policy or job. |
| **Event age** | Displays in numeric form how long ago from the current time the event occurred. |

You can click View All at the bottom of the event list to navigate the browser to the Notifications page.

## Clusters map

The **Clusters map** indicates the geolocation of each cluster, using red, orange, and green markers on the map.

The colored markers indicate the following:

- **Red**: At least one required service has stopped on the cluster.
- **Orange**: All required services are running but the remaining capacity on a cluster is less than 10%.
- **Green**: All required services are running and remaining disk capacity is greater than 10%.

You can move the cursor over a marker on the map displays a tooltip specifying the data center associated with the cluster, the cluster name, and the number of DLM policies that are associated with that cluster.

You can click a marker to open a panel showing the same information as in the tooltip, plus a Launch Ambari link. Clicking the link opens a new browser tab with the login page for the Ambari host for that cluster.

If the dot is red, the panel also displays a list of services that are in a Stopped state in Ambari.

## Issues & Updates table

The **Issues & Updates table** shows policies that have running jobs but at least one failed out of the most recent 10 jobs. You do not see any policy if its last 10 jobs were all successful.

Table columns include the following:

| | |
|---|---|
| **Job Status** | When the status of a job is Running, a status circle icon and progress bar display. For jobs that are not running, a status circle icon displays along with the text Success, Failed, or Ignored. You can move the cursor over a |

Failed status to see a “View Log” tooltip, which you can click to see the job log.

**Source & Destination** The names of the source and destination clusters associated with the policy.

**Service** Indicates whether the data being replicated is from HDFS or Hive.

**Policy** The name assigned to the policy.

**Policy History** Shows up to 10 job statuses as colored dots.

| Color | Status | Description |
|---|---|---|
| Green | Succeeded | Job completed with no issues. |
| Red | Failed | Job did not complete. |
| Gray | Ignored | Job did not start because a previous ins Only one run of a job can be in progres ignored, you might need to modify its f |

Clicking the colored dots navigates the browser to the Policies page, with the filter preset to show information only about the specified policy.

**Transferred/Files** The amount of data transferred, in gigabytes, and the number of objects transferred, if available. When a job is running, the column displays In Progress.

**Runtime** How long it took to run the most recent job.

**Started** When the most recent job started.

**Ended** When the most recent job ended.

**Actions icon**

- **Abort Job**: Aborts a running job. Enabled only when the job status is Running.
- **Re-run Job**: Starts another instance of the policy. Disabled when a job is executing.
- **Edit Policy**: Allows editing of some policy settings. Disabled if a policy is expired.
- **Delete Policy**: Removes a policy from Data Lifecycle Manager. Delete cannot be undone. Always enabled.
- **Suspend Policy**: Suspends the policy and any job that is executing. Disabled when the policy status is Suspended.
- **Activate Policy**: Resumes a suspended policy. Disabled when the policy status is Running.

# Preparing to setup replication policy

You must understand the various roles and permissions that are required to work with data replication policies.

## Roles required

DLM has the following roles for the users:

### Infrastructure Admin role

The **Infrastructure Administrator** has access to DLM and administrative permissions to perform all actions in DLM.

- Can access DP Platform service to manage clusters enabled for Data Lifecycle Manager. Addition of clusters also requires **Ambari admin** or **Cluster admin** privileges in the Ambari service of that HDP cluster.

  For more information, see Apache Ambari Administration.
- Can use all DLM functionality, including enabling snapshot based replication

### DLM Admin

The **DLM Administrator** can perform most of the operations in the DLM just as **Infra administrator**.

The DLM admin has the following capabilities and restrictions:

- Can pair Hadoop clusters.
- Can add, modify, and delete cloud credentials.
- Can create, edit, suspend, activate, and delete replication policies.

  

  **Note:** To create and manage replication policies, the DLM admin must at least have the **Cluster User** role assigned for the corresponding source and destination clusters in Ambari. For more information, see Understanding cluster roles and access documentation for Ambari administration.

  - Cannot setup snapshot-based replication policies.

### DLM User

The **DLM user** has the following capabilities and restrictions:

- View cluster listing and pairing information
- View the DLM policy listing page
- Can monitor replication policies
- Can abort the already running policy
- Can rerun any failed policy
- Cannot pair clusters
- Cannot modify a replication policy

## Working with clusters

Some of the operations that you can perform while working with clusters.

### Add clusters

Any source or destination cluster that you want to use in a replication relationship with DLM must be managed by Apache Ambari and enabled for DLM through DPS Platform.

You must have a cluster registered with the **Data Lifecycle Manager** to which you replicate data. The cluster must have enough storage to accept data that gets replicated.

#### Cluster pairing

Before setting up replication between two hadoop clusters, the clusters need to be paired which will validate that the data and metadata can be replicated between two clusters. This will verify the configurations of two clusters to communicate with each other.

#### Procedure

1. In the DLM navigation pane, click **Pairings**.
2. Click **Add Pairing**.

   The Create Pairing page displays, showing the clusters that are enabled for replication.

**Tip:** You can place the cursor over the cluster name to display the cluster location.

3. Click one of the cluster names.

   All clusters available to be paired with the cluster you selected display in a second column.

**Tip:** If a cluster displays but cannot be selected, it is already paired with the cluster you selected in the first column.

4. Click a cluster in the second column.
5. Click **Start Pairing**.

   A progress bar displays.
6. Repeat the above steps to pair additional clusters.

#### Pairing considerations

You should take into consideration the following items when pairing clusters in DLM.

- For pairing to succeed, host name resolution must work between all the nodes involved (DPS Platform host and all cluster hosts.)

  For example, pairing in DLM fails if the DLM engines on the clusters being paired cannot resolve each other's host name.
- You can only pair clusters that have been registered with DPS Platform and enabled for use with Data Lifecycle Manager.
- The HDFS nameservice for the source and destination clusters cannot be configured with the same name.
- Cluster security configurations must be symmetrical to pair clusters, including LDAP, Kerberos, Ranger, and Knox.

# Cloud credentials

## Register cloud credentials

If you plan to replicate data to a storage cloud account, you must register the cloud credentials, so DLM can access your cloud account. There are three types of cloud storage account that are supported; **Amazon S3**, **Microsoft WASB**, and **Google Cloud Storage**.

### Registering Amazon S3 cloud account

You must have valid **Amazon S3** credentials to register the cloud account with DLM.

#### Procedure

1. In the DLM UI navigation pane, click **Cloud Credentials > Add**
2. Enter the details in the **Add Cloud Credential** window:
   - **Cloud Storage Type** - Select the replication cloud account from the drop-down.
   - **Name** - Provide a unique cloud credential name.
   - **Authentication Type** - Select the authentication type as **Access Secret Key** from the drop-down.

**Note:** If you select **IAM Role**, click **Save** to proceed.

   - **Access Key** - Enter the valid access key.
   - **Secret Key** - Enter the valid secret key.

3. Click **Validate**.

   Using the validation feature is recommended to ensure that the Amazon S3 bucket keys are valid. If the keys are not valid, the DLM policy cannot execute a copy of data to the target Amazon S3 bucket.

   Verify that your credentials are listed on the **Cloud Credentials** page.

#### Considerations for Amazon S3

Cloud bucket requirements

- You need a cloud bucket with user credentials that you can enter in DLM, so DLM can access the bucket.
- The bucket has to have enough space for the replicated data, and write permissions to copy the data.
- The bucket needs to support cloud storage encryption types supported by DLM (SSE-S3 & SSE-KMS).

## Registering Microsoft WASB cloud account

You must have valid **WASB** credentials to register the cloud account with DLM.

#### Procedure

1. Create a storage account in WASB. The **Access key** can be retrieved from the WASB storage account. The Access Key is used in DLM UI to set up the cloud account credentials.
2. In the DLM UI navigation pane, select **Cloud Credentials > Add**.
3. Enter the WASB cloud details in the **Add Cloud Credential** window:
   - **Cloud Storage Type** - Select the replication cloud account from the drop-down.
   - **Name** - Provide a unique cloud credential name.
   - **Storage Account Name** - Provide a name for the storage account.
   - **Access Key** - Paste the Access Key generated from the newly created storage account.
4. Click **Save** to save the changes. Verify that your credentials are listed on the **Cloud Credentials** page.

#### Considerations for Microsoft WASB

Requirements for Blob containers

- You need a cloud account with user credentials that you can enter in DLM, so DLM can access Blob containers.
- Create a WASB storage account using: https://portal.azure.com
- Blob containers must have enough space for the replicated data, and write permissions to copy the data. For more information about WASB, see https://blogs.msdn.microsoft.com/cindygross/2015/02/04/understanding-wasb-and-hadoop-storage-in-azure/

## Registering Google cloud account

You must have valid **Google cloud** credentials to register the cloud account with DLM.

#### Procedure

1. Create a storage account in Google Cloud Storage. The **Access & Secret key** can be retrieved from the Google storage account. This key is used in DLM UI to set up the cloud account credentials.
2. In the DLM UI navigation pane, select **Cloud Credentials > Add**.
3. Enter the Google cloud details in the **Add Cloud Credential** window:
   - **Cloud Storage Type** - Select the replication cloud account from the drop-down.
   - **Name** - Provide a unique cloud credential name.
   - **UPLOAD File** - You must upload the service account private key that you received when you created the Google cloud storage account.

     

     **Note:** This file must be available in JSON format only.
4. Click **Save** to save the changes. Verify that your credentials are listed on the Cloud Credentials page.

#### Considerations for Google Cloud Storage

Cloud bucket requirements

- You need a cloud bucket with user credentials that you can enter in DLM, so DLM can access the Google cloud bucket.
- The bucket has to have enough space for the replicated data, and write permissions to copy the data.

# Data replication use cases

This page provides information regarding the data replication use cases pertaining to **HDFS** and **HIVE** services.

## Replication of HDFS data

HDFS data can be replicated using multiple entities

The DLM App submits the replication policy to the DLM Engine on the destination cluster. The DLM Engine then schedules replication jobs at the specified frequency.

- At the specific frequency, DLM Engine submits a DistCp job that runs on destination YARN, reads data from source HDFS, and writes to destination HDFS.
- File length and checksums are used to determine changed files and validate that the data is copied correctly.
- The Ranger policies for the HDFS directory are exported from source Ranger service and replicated to destination Ranger service.
- Atlas entities that related to HDFS directory are replicated. If no HDFS path entities are present within Atlas, they are created and then exported.
- Atlas replication is optional during the DLM policy creation.
- DLM Engine also adds a deny policy on the destination Ranger service for the target directory so that the target is not writable.

### HDFS on-premise replication

Before you can begin replicating data using clusters, you must make sure that there are at least a couple of clusters that are registered in your DLM App instance.

#### Replication of data on-premise to on-premise in HDFS

You must create a replication policy that specifies the data to replicate, the replication schedule, and other settings.

#### About this task

You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

#### Procedure

1. Pair clusters for replication. Select the two clusters to use for replication and pair them, so they can communicate with each other. For more information, see Pair clusters for replication.
2. Create a replication policy.
3. Select **Policies** and click **Add Policy**. By default, **HDFS** is selected as the service in the **Create Replication Policy** page.
4. Enter the replication policy name and description.
5. Click **SELECT SOURCE** and select type and source cluster from the drop-down.
6. Provide the data replication folder path and click **SELECT DESTINATION**.
7. Select the destination type from the drop-down.

   You must select another cluster available in the DLM App as your destination.

8. Select the path and click **VALIDATE**.
9. Once the validation is successful, click **SCHEDULE**.
10. Configure the job settings for the replication policy.
11. Click **ADVANCED SETTINGS** to set up the policy queue.
12. Click **CREATE POLICY**.

    The data replication process is enabled.

    View job status from the policies page. Verify that the job starts and runs as expected.

## HDFS cloud replication

DLM supports replication of HDFS data from cluster to cloud storage and vice versa. The replication policy runs on the cluster and either pushes or pulls the data from cloud storage.

The cluster can be an On-premise or IaaS cluster with data on local HDFS. The cluster requires HDFS, YARN, Ranger, Knox and DLM Engine services.

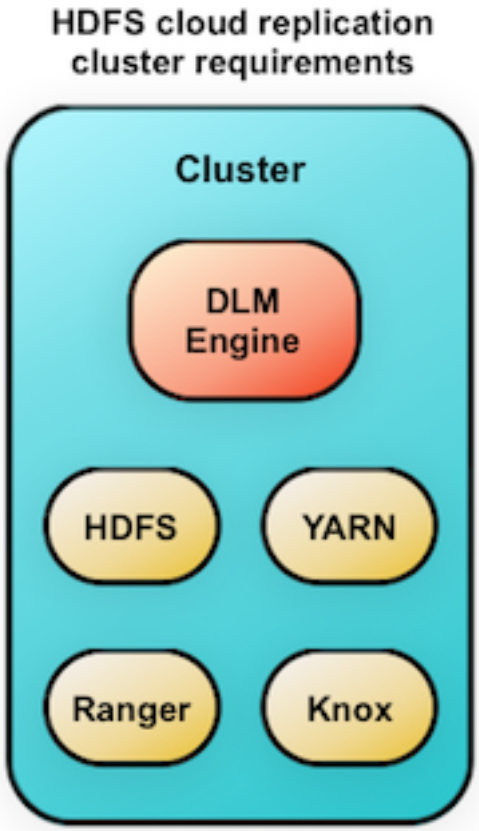


The HDFS replication policy process overview consists of the following:

- The DLM App submits the replication policy to the DLM Engine on the destination cluster. The DLM Engine then schedules replication jobs at the specified frequency.
- At the specific frequency, DLM Engine submits a DistCp job that runs on destination YARN, reads data from source HDFS, and writes to destination HDFS.
- File length and checksums are used to determine changed files and validate that the data is copied correctly.
- The Ranger policies for the HDFS directory are exported from source Ranger service and replicated to destination Ranger service.

**Note:** DLM Engine also adds a deny policy on the destination Ranger service for the target directory so that the target is not writable.

- Atlas entities related to HDFS directory are replicated. If there are no HDFS path entities are present within Atlas, they are first created and then exported.

### On-premise to Amazon S3 replication in HDFS

The process for creating a replication job from on-premise to Amazon S3 is similar to creating one for on-premise to on-premise. The primary difference is that, you must register your cloud account credentials with DLM App instance, so that DLM can access your cloud storage.

Attention: Replication of HDFS data from on-premise to cloud is a limited GA feature in DPS 1.1. The HDFS data that you replicate to cloud requires security policies outside the Hadoop system, so you should work with Hortonworks support to ensure proper configuration of your environment. This does not apply to Hive replication to cloud.

### Replication of data on-premise to Amazon S3 in HDFS

You must create a new replication policy to replicate data from on-premise Amazon S3 cloud storage.

#### About this task

Before you create a new replication policy, you must register Amazon S3 cloud account. For more information, see Register cloud credentials. You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

**Note:** You can replicate data on-premise to Amazon S3 with a single cluster.

#### Procedure

1. Select **Policies** and click **Add Policy**. By default, **HDFS** is selected as the service in the **Create Replication Policy** page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE** and select type and source cluster from the drop-down.
4. Provide the data replication folder path and click **SELECT DESTINATION**.
5. Select the destination type as **S3** and **Cloud Credential** from the drop-down.
6. Provide a folder path bucket_name/path and click **VALIDATE**.
7. Once the validation is successful, click **SCHEDULE**.
8. Configure the job settings for the replication policy.
9. Click **ADVANCED SETTINGS** to set up the policy queue.
10. Click **CREATE POLICY**.

    The data replication process is enabled.

    View job status from the policies page. Verify that the job starts and runs as expected.

### Amazon S3 to on-premise replication in HDFS

You must have a cloud account registered in Amazon S3 before you perform data replication from Amazon S3 to on-premise.

### Replication of data from Amazon S3 to on-premise in HDFS

You must create a new replication policy to replicate data from Amazon S3 cloud storage to on-premise.

#### About this task

Before you create a new replication policy, you must register Amazon S3 cloud account. For more information, see Register cloud credentials. You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

### Procedure

1. Select **Policies** and click **Add Policy**. By default, **HDFS** is selected as the service in the **Create Replication Policy** page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE**.
4. Select type as **S3** and **Cloud Credential** from the drop-down and enter the S3 source path bucket_name/path.
5. Click **SELECT DESTINATION**.

   Make sure you have one or more clusters in the DLM application.
6. Select type as cluster and destination cluster from the drop-down.
7. Enter the destination path and click **VALIDATE**.
8. Once the validation is successful, click **SCHEDULE**.
9. Configure the job settings for the replication policy.
10. Click **ADVANCED SETTINGS** to set up the policy queue.
11. Click **CREATE POLICY**.

    The data replication process is enabled.

    View job status from the policies page. Verify that the job starts and runs as expected.

### On-premise to Microsoft WASB replication in HDFS

The process for creating a data replication job from on-premise to WASB is similar to creating one for on-premise to on-premise. The primary difference is that, you must register your WASB cloud credentials with DLM App instance, so that DLM can access your WASB cloud storage. You must create a new data replication policy to replicate data from on-premise to WASB.

### Replication of data on-premise to Microsoft WASB in HDFS

You must create a new replication policy to replicate data from on-premise to WASB cloud account.

### About this task

Before you create a new replication policy, you must register the WASB cloud account. For more information, see Register cloud credentials. You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

### Procedure

1. Select **Policies** and click **Add Policy**. By default, **HDFS** is selected as the service in the **Create Replication Policy** page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE** and select type and source cluster from the drop-down.
4. Provide the data replication folder path and click **SELECT DESTINATION**.
5. Select the destination type as **WASB** and **Cloud Credential** from the drop-down.
6. Provide a folder path container_name/path and click **VALIDATE**.
7. Once the validation is successful, click **SCHEDULE**.
8. Configure the job settings for the replication policy.
9. Click **ADVANCED SETTINGS** to set up the policy queue.
10. Click **CREATE POLICY**.

    The data replication process is enabled. View job status from the policies page. Verify that the job starts and runs as expected.

### Microsoft WASB to on-premise replication in HDFS

You must setup cloud storage account in WASB before you perform replication from WASB cloud storage to on-premise. Later, create a new replication policy to replicate data from WASB to on-premise.

### Replication of data from Microsoft WASB to on-premise in HDFS

You must create a new replication policy to replicate data from WASB cloud account to on-premise.

#### About this task

Before you create a new replication policy, you must register the WASB cloud account. For more information, see Register cloud credentials. You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

**Note:** You must have a cluster registered with the Data Lifecycle Manager to which you replicate data. The cluster must have enough storage to accept data that gets replicated.

#### Procedure

1. Select **Policies** and click **Add Policy**. By default, **HDFS** is selected as the service in the **Create Replication Policy** page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE**.
4. Select type as **WASB** and **Cloud Credential** from the drop-down and enter the path container_name/path for the WASB source.
5. Click **SELECT DESTINATION**.

   You must have one or more clusters in the DLM application.
6. Select cluster type and destination cluster from the drop-down.
7. Enter the destination path and click **VALIDATE**.
8. Once the validation is successful, click **SCHEDULE**.
9. Configure the job settings for the replication policy.
10. Click **ADVANCED SETTINGS** to set up the policy queue.
11. Click **CREATE POLICY**.

    The data replication process is enabled. View job status from the policies page. Verify that the job starts and runs as expected. For more information, see Viewing job status.

### On-premise to Google Cloud replication in HDFS

The process for creating a replication job from on-premise to Google Cloud Storage is similar to creating one for on-premise to on-premise. The primary difference is that, you must register your Google cloud account credentials with DLM App instance, so that DLM can access your cloud storage.

### Replication of data from on-premise to Google Cloud Storage in HDFS

You must have a cluster registered with the DLM app to perform data replication from on-premise to Google cloud. You must register your cloud credentials. For more information, see Register cloud credentials.

#### About this task

You must create a new replication policy to replicate data from on-premise to Google cloud account. You can replicate data on-premise to Google cloud storage using a single cluster. You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

#### Procedure

1. Select **Policies** and click **Add Policy**.
   By default, **HDFS** is selected as the service in the Create Replication Policy page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE** and select type and source cluster from the drop-down.
4. Provide the data replication folder path and click SELECT DESTINATION.
5. Select the destination type as **GCS** and **Cloud Credential** from the drop-down.
6. Provide a folder path bucket_name/path and click **VALIDATE**.

7. Once the validation is successful, click **SCHEDULE**.
8. Configure the job settings for the replication policy.
9. Click **ADVANCED SETTINGS** to set up the policy queue.
10. Click **CREATE POLICY**.
    The data replication process is enabled.

    View job status from the policies page. Verify that the job starts and runs as expected.

#### Google Cloud to on-premise replication in HDFS

You must setup a cloud storage account in Google Cloud Storage before you perform replication from Google cloud storage to on-premise.

#### Replication of data from Google Cloud Storage to on-premise in HDFS

You must have a cluster registered with the Data Lifecycle Manager to which you want to replicate data from Google Cloud. The cluster must have enough storage to accept data that gets replicated. For more information, see Register cloud credentials.

#### About this task

You must create a new replication policy to replicate data from Google cloud storage to on-premise. You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

#### Procedure

1. Select **Policies** and click **Add Policy**.
   By default, **HDFS** is selected as the service in the **Create Replication Policy** page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE**.
4. Select type as **GCS** and **Cloud Credential** from the drop-down and enter the path container_name/path for the GCS source.
5. Click **SELECT DESTINATION**.

   You must have one or more clusters in the DLM application.
6. Select cluster type and destination cluster from the drop-down.
7. Enter the destination path and click **VALIDATE**.
8. Once the validation is successful, click **SCHEDULE**.
9. Configure the job settings for the replication policy.
10. Click **ADVANCED SETTINGS** to set up the policy queue.
11. Click **CREATE POLICY**.
    The data replication process is enabled.

    View job status from the policies page. Verify that the job starts and runs as expected.

# Replication of HIVE data

Hive data replication involves multiple entities

## Hive replication concepts

DLM supports Hive replication.

#### Hive tables - Managed and External

Managed tables are Hive owned tables where the entire lifecycle of the tables' data are managed and controlled by Hive. External tables are tables where Hive has loose coupling with the data.

All the write operations to the managed tables are performed using Hive SQL commands. If a managed table or partition is dropped, the data and metadata associated with that table or partition are deleted. The transactional semantics (ACID) are also supported only on managed tables.

The writes on external tables can be performed using Hive SQL commands but data files can also be accessed and managed by processes outside of Hive. If an external table or partition is dropped, only the metadata associated with the table or partition is deleted but the underlying data files stay intact. A typical example for external table is to run analytical queries on HBase or Druid owned data via Hive, where data files are written by HBase or Druid and Hive reads them for analytics.

Hive supports replication of external tables with data to target cluster and it retains all the properties of external tables.

The data files permission and ownership are preserved so that the relevant external processes can continue to write in it even after failover.

For handling conflicts in external tables' data location due to replication from multiple source clusters to same target cluster, DLM assigns a unique base directory for each source cluster under which, external tables data from corresponding source cluster would be copied. For example, if external table location at a source cluster is /ext/hbase_data and after replication, the location in target cluster would be <base_dir>/ext/hbase_data. Users can track the new location of external tables using DESCRIBE TABLE command.

**Attention: DLM upgrade use-case**: In a normal scenario, if you had external tables that were replicated as managed tables, after the upgrade process, you must drop those tables from target and set the base directory. In the next instance they get replicated as external tables.

### Bootstrap and incremental replication

DLM allows you to replicate Hive databases from a source cluster to a target location on a destination cluster.

When you initiate the replication of Hive data, all of the data from the source location is copied to the destination. This bootstrapping of data can take hours to days, depending on factors such as the amount of data being copied and available network bandwidth. Subsequent replication jobs from the same source location to the same target on the destination are incremental, so only the changed data is copied.

If a bootstrap replication is interrupted, such as due to a network failure or an unrecoverable error, DLM automatically retries the job. If a retry succeeds, the replication job continues from the point at which it was interrupted. If the automatic retries are not successful, you must manually correct the problem before running the policy again. When you activate the policy again, the replication job resumes from the point at which it was suspended.

After the bootstrap replication succeeds, an incremental replication is automatically performed. This job synchronizes, between the source and destination clusters, any events that occurred during the bootstrap process. After the data is synchronized, the replicated data is ready for use on the destination.

Functions such as User Defined Functions (UDF) in Hive are replicated. To enable this, UDFs have to be created using a syntax. An example of UDF creation syntax:

```
CREATE FUNCTION [db_name.]function_name AS class_name  USING JAR|FILE|
ARCHIVE 'file_uri' [, JAR|FILE|ARCHIVE 'file_uri'] ;
```

- ACID tables, external tables, storage handler-based tables (such as HBase), and column statistics are currently not replicated.
- When creating a schedule for a Hive replication policy, you should set the frequency so that changes are replicated often enough to avoid overly large copies.

**Incremental Replication**

The incremental replication in Hive is achieved using notification events maintained by Hive in Hive Metastore.

Hive logs notification events for all operations (both metadata and data changes) on the managed table but in case of external tables, data writes cannot be tracked by Hive as it is performed by external sources directly without using

Hive SQL commands. Therefore, Hive always copies the latest data from external tables to target cluster to avoid any loss of data.

### Storage-based authorization

Hive supports doAs=true plus storage-based authorization that enables security at Hive Metastore API level.

This mode does not involve any additional enforcement at SQL level (in HiveServer2). Customer applications often use this mode when they do not require fine grained access control at column or row level. In this mode, the files are typically owned by the end user. The queries run as end user using doAs=true, and permissions are provided to end-user to access the files directly, since HDFS permissions are set appropriately.

When Hive replication is performed with this mode, the file permissions need to be preserved in target cluster as well, so that the end user shall continue to access the replicated data files based on appropriate permissions.

### Statistics replication

Basic statistics such as the number of rows of a table or partition and the column statistics such as histograms (min, max, count) of a particular interesting column are important in many ways.

One of the key use cases of statistics is query optimization.

Hive supports Cost Based Optimizer (CBO) which primarily depends on column statistics to optimize the query execution plan. For example, if the min and max values of a integer type column “c” in a partition “p” is min=10 and max=50, then a query with predicate such as c < 10 or c > 50 or anything that does not fall under the range, shall avoid scanning the partition “p”. For transactional tables, column statistics are supported only if hive.txn.stats.enabled=true.

Currently, there are two modes to compute statistics:

- ANALYZE command: analyze table t [partition p] compute statistics for [columns c,...];
- Hive automatically computes the statistics whenhive.stats.autogather=true for basic statistics andhive.stats.column.autogather=truefor column statistics.

As statistics are key for query optimization, when table/partition data is replicated to target cluster, it is important to replicate the statistics as well which would speed up the queries running on target clusters. One of the key requirement for statistics replication to work accurately is it’s consistency with current dataset in a table or partition. Hive replication can achieve it with point-in time consistent incremental replication model. If any database/table is bootstrapped, then corresponding basic and columns statistics (if present) would also be bootstrapped.

As Hive replication takes care of replicating statistics to target cluster, it is recommended to disable hive.stats.autogatherandhive.stats.column.autogather at the target cluster. If it is enabled, it would cause additional computation cost in the target cluster.

**Important:** For existing policies, any statistics gathered after deploying DLM 1.4.0 will be replicated to the target cluster. For any new policies created using DLM 1.4.0 release onwards, statistics is replicated to the target cluster.

### Replication differences between HDP 2.6.5 to 3.x

In HDP 2.6.5 version, managed tables are managed by Hive.

In HDP 2.6.5 version, managed tables are operated by Hive. The managed tables can be any file format like ORC, AVRO or Flat file, and can also be both transactional and non-transactional. But in HDP 3.x version, managed tables are always transactional (either full ACID or micro-managed tables).

When Hive tables are replicated from HDP 2.6.5 to 3.x, the source tables get converted to transactional tables or external tables at target, and remain so throughout the replication process.

The conversion criteria is:

- Managed tables with ORC file format converted to full ACID tables.
- Managed tables with file format other than ORC will be converted to transactional tables having **Insert-Only** support.

- Managed tables with external AVRO schema or multi-level directories or non-native tables will be converted as external tables.

### Hive on-premise replication

Before you can begin replicating data using clusters on Hive, you must make sure that there are at least a couple of clusters that are registered in your DLM App instance. The replication load happens on the target cluster.

### Replication of data on-premise to on-premise in Hive

You must create a replication policy that specifies the data to replicate, the replication schedule, and other settings.

#### About this task

You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

#### Procedure

1. Select **Policies** and click **Add Policy**. Select **HIVE** as the service in the **Create Replication Policy** page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE** and select type and source cluster from the drop-down.
4. Provide the data replication folder path and click **SELECT DESTINATION**.
5. Select the destination type from the drop-down.

   You must select another cluster available in the DLM App as your destination.
6. Select the path and click **VALIDATE**.
7. Once the validation is successful, click **SCHEDULE**.
8. Configure the job settings for the replication policy.
9. Click **ADVANCED SETTINGS** to set up the policy queue.
10. Click **CREATE POLICY**.

    The data replication process is enabled.

    View job status from the policies page.

    Verify that the job starts and runs as expected.

## Hive cloud replication

DLM supports replication of the Hive database from a cluster with underlying HDFS to another cluster with cloud storage. It uses push-based replication, with the replication job running on the cluster with HDFS.

Hive stores its metadata in Hive Metastore, but the underlying data is stored in HDFS or cloud storage. In a Hadoop cluster with Hive service, the Hive warehouse directory can be configured with either HDFS or cloud storage.

**Note:** Hive replication from cloud storage to HDFS is currently not supported.

- You can rename the dataset in the policy that is replicated.
- You can create a pull-based policy on the source cluster to move data from the target back into the source cluster Hive database.
- DLM does not manage Ranger policies and any PII/secure data that gets replicated from on-premise to Amazon S3. You must manage these items outside of DLM.
- Hive replication from an HDFS-based cluster to a cloud storage-based cluster requires the following:
  - Source cluster

    The cluster with a Hive warehouse directory on local HDFS. This can be an On-premise cluster or an IaaS cluster with data on local HDFS. The required services are HDFS, YARN, Hive, Ranger, Knox and DLM Engine.
  - Destination cluster

The cluster with data on cloud storage. The cluster minimally requires Hive Metastore, Ranger, Knox, and DLM Engine.

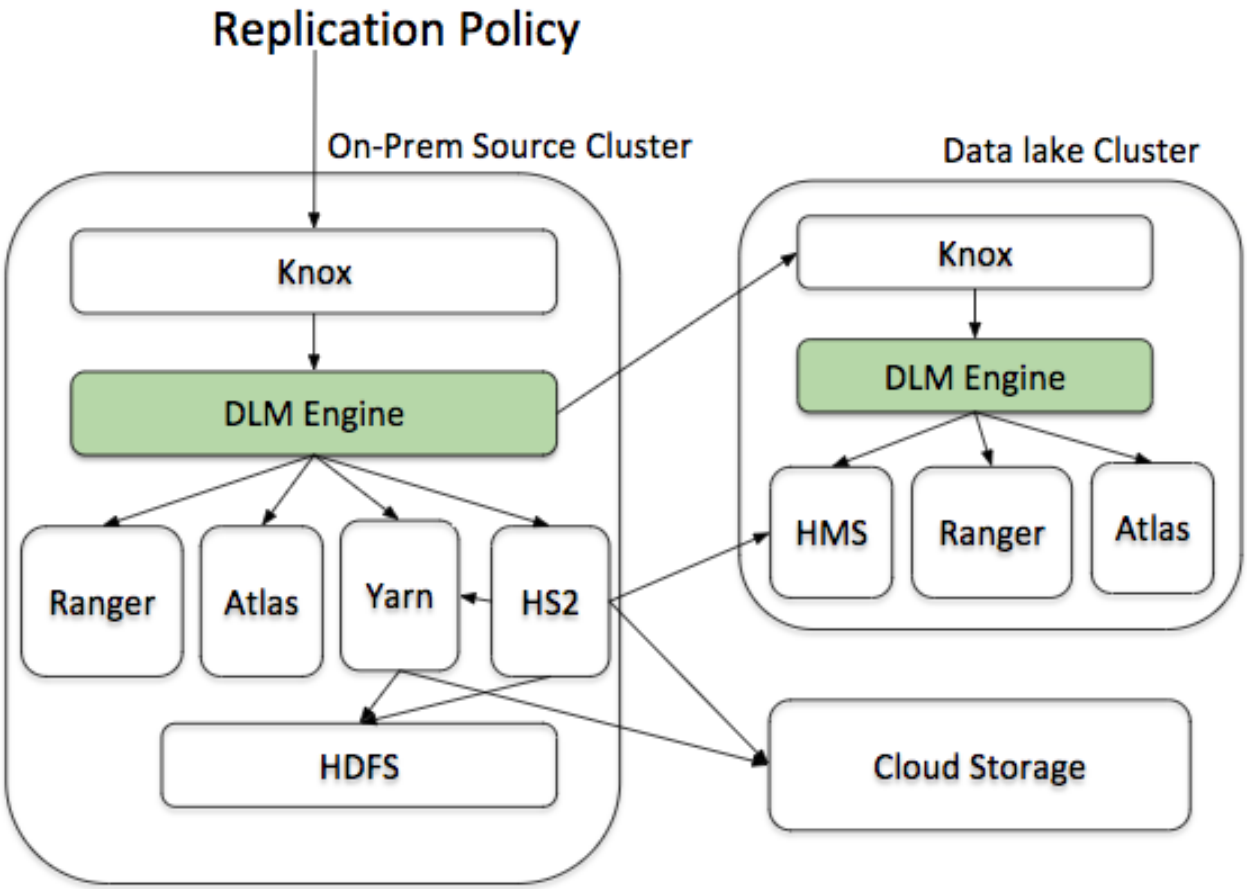


## Setting target cluster for cloud storage in Hive

Before performing Hive replication from on-prem to any supported cloud storage, the target cluster for Hive cloud replication should be set up on cloud storage instances, with Hive warehouse directory on that specific cloud storage.

The target cluster is Data Lake cluster with metadata services such as HMS, Ranger, Atlas, and DLM Engine.

For a specific cloud account that is used for data replication, you must set up applicable path values for Hive replication function and Hive metastore parameters.

### Amazon S3 cloud storage

When you set up Amazon S3 as your target cloud cluster, use the following Hive metastore configuration:

```
hive.metastore.warehouse.dir=s3a://<bucket_name>/<warehouse_path>
```

The target cluster must have additional Amazon S3 credential configurations to access Amazon S3 storage buckets. For more information, see Configuring Access to S3.

### Microsoft WASB cloud storage

When you set up WASB as your target cloud cluster, use the following Hive metastore configuration:

```
hive.metastore.warehouse.dir=wasb://
<container_name>@<storage_account_name>.blob.core.windows.net/
<warehouse_path>
```

The target cluster must have additional WASB credential configurations to access WASB storage containers. For more information, see Configuring Access to WASB.

### Google cloud storage

When you set up Google cloud as your target cloud cluster, use the following Hive metastore configuration:

```
hive.metastore.warehouse.dir=gs://<bucket_name>/<warehouse_path>
```

The target cluster must have additional Google cloud storage credential configurations to access Google cloud buckets.

Add and save the following configurations in core-site.xmlfile.

fs.gs.auth.service.account.email=email id of gcs service account

fs.gs.auth.service.account.private.key.id=private key id of gcs service account

fs.gs.auth.service.account.private.key=private key of gcs service account

The values for these configurations can be found in the JSON file that you downloaded while registering the Google cloud storage credentials with the DLM App.

**Note:** For more information, see Registering Google Cloud Account.

### On-premise to Amazon S3 replication in HIVE

The process for creating a Hive data replication job from on-premise to Amazon S3 is similar to creating one for on-premise to on-premise. The primary difference is that, you must register your cloud account credentials with DLM App instance, so that DLM can access your cloud storage. The replication load happens on the source cluster.

### Replication of data on-premise to Amazon S3 in Hive

You must create a new data replication policy to replicate data from on-premise to Amazon S3. You must setup target cluster before commencing the replication process.

### About this task

Before you create a new replication policy, you must register Amazon S3 cloud account with the DLP App. For more information, see Register cloud credentials. You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

**Note:** You can replicate data on-premise to Amazon S3 with a single cluster. The metastore must be running on the cloud. There is no requirement to run the HiveServer 2 on the cloud environment.

### Procedure

1. Select **Policies** and click **Add Policy**. Select **HIVE** as the service in the **Create Replication Policy** page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE** and choose **Type**, **Source Cluster**, and **Select Database**.
4. Click **SELECT DESTINATION** and choose **Type** and **Destination Cluster**.
5. Enter the **Destination Database**.
6. Provide the **Hive External Table Base Directory** path: S3://bucket_name/path

The external table base directory path cannot be changed once the policy is created.

7. Select **Cloud Credential** from the drop-down.
8. Click **VALIDATE**.
9. Once the validation is successful, click **SCHEDULE**.
10. Configure the job settings for the replication policy.
11. Click **ADVANCED SETTINGS** to set up the policy queue.
12. Click **CREATE POLICY**.

The data replication process is enabled.

View job status from the policies page. Verify that the job starts and runs as expected.

### On-premise to Microsoft WASB replication in HIVE

The process for creating a Hive data replication job from on-premise to WASB is similar to creating one for on-premise to on-premise. The primary difference is that, you must register your WASB cloud credentials with DLM App instance, so that DLM can access your WASB cloud storage.

### Replication of data on-premise to Microsoft WASB in Hive

You must create a new data replication policy to replicate data from on-premise to WASB. You must setup target cluster before commencing the replication process.

#### Before you begin

You must register the WASB cloud account with the DLP App. For more information, see Register cloud credentials.
You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

#### About this task

You can replicate data on-premise to WASB with a single cluster. The metastore must be running on the cloud. There is no requirement to run the HiveServer 2 on the cloud environment.

#### Procedure

1. Select **Policies** and click **Add Policy**. Select **HIVE** as the service in the **Create Replication Policy** page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE** and choose **Type**, **Source Cluster**, and **Select Database**.
4. Click **SELECT DESTINATION** and choose **Type** and **Destination Cluster**.
5. Enter the **Destination Database**.
6. Provide the **Hive External Table Base Directory** path: WASB://container_name/path

   The external table base directory path cannot be changed once the policy is created.
7. Select **Cloud Credential** from the drop-down.
8. Click **VALIDATE**.
9. Once the validation is successful, click **SCHEDULE**.
10. Configure the job settings for the replication policy.
11. Click **ADVANCED SETTINGS** to set up the policy queue.
12. Click **CREATE POLICY**.

The data replication process is enabled.

View job status from the policies page. Verify that the job starts and runs as expected.

### On-premise to Google Cloud replication in HIVE

The process for creating a Hive data replication job from on-premise to Google cloud storage is similar to creating one for on-premise to on-premise. You must register your GCS cloud credentials with DLM App instance, so that DLM can access your Google cloud storage.

### Replication of data on-premise to Google Cloud in HIVE

#### Before you begin

You must register the Google Cloud Storage cloud account with the DLP App. For more information, see Register cloud credentials.

#### About this task

You can replicate data on-premise to Google cloud with a single cluster. The metastore must be running on the cloud. There is no requirement to run the HiveServer 2 on the cloud environment. You must have **Infra Admin** or **DLM Admin** role to perform this set of tasks.

#### Procedure

1. Select **Policies** and click **Add Policy**. Select **HIVE** as the service in the **Create Replication Policy** page.
2. Enter the replication policy name and description.
3. Click **SELECT SOURCE** and choose **Type**, **Source Cluster**, and **Select Database**.
4. Click **SELECT DESTINATION** and choose **Type** and **Destination Cluster**.
5. Enter the **Destination Database**.
6. Provide the **Hive External Table Base Directory** path: GCS://bucket_name/path

   The external table base directory path cannot be changed once the policy is created.
7. Select **Cloud Credential** from the drop-down.
8. Click **VALIDATE**.
9. Once the validation is successful, click **SCHEDULE**.
10. Configure the job settings for the replication policy.
11. Click **ADVANCED SETTINGS** to set up the policy queue.
12. Click **CREATE POLICY**.

    The data replication process is enabled.

    View job status from the policies page. Verify that the job starts and runs as expected.

# Metadata replication

This page provides information about various types of metadata replication.

## Ranger metadata

When a DLM replication job is run, data, metadata, and any Ranger policies that are associated with the replicated data are automatically exported to the target.

The replicated data on the destination is marked as read-only by adding a deny policy on the replicated data in Ranger in the destination cluster. This prevents accidental writes on the copy.

For on-premise to on-premise replications, the policies, permissions, and ACLs are retained and applied to the data on the target, except that the destination data is read-only.

For on-premise to cloud replication, the Ranger policies, permissions, and ACLs are stored in metadata files in cloud storage. Data in the cloud is protected using security features in the cloud environment.

## Atlas metadata

Atlas plugin within DLM is used to replicate Atlas metadata. It uses incremental export to move data across clusters, thereby optimizing the payload for speed and size. Atlas replication is optional and can be turned ON during the DLM replication policy setup time.

Atlas entities replicated to target cluster are tagged with special classification. Tagging entities allows for easy access to the entities that are part of the available metadata due to replication.

**Note:** The lineage associated with the entities is not replicated.

On the source cluster, the entity's replicatedTo attribute is updated to indicate the cluster it is being replicated to and on the target cluster the entity's replicatedFrom attribute is modified to indicate its source. Since each cluster has its own identity, the entities that are part of replication are transformed such that, they appear to be native to the cluster they are going to reside within. This involves changing attributes that are indicative of their place of residence. In addition, within Atlas, new entities of type AtlasServer are created. This allows for a central place to access all the servers for which replication has been initiated. Replication audit logs can also be accessed here. Each audit entry has details of every export or import performed for that cluster.

When a DLM Atlas replication job is executed, any Atlas metadata associated with the dataset on source Atlas server, which is replicated, is exported from source, and imported in the target Atlas cluster. The associated replication policy must not be updated or modified during the course of the replication life cycle. You can perform Atlas replication on-premise to on-premise using both HDFS and Hive. You must make sure that there are at least two clusters that are registered in your DLM app instance. And Atlas must be installed on source and target clusters. Optionally, using Ambari UI, you can verify if Atlas is installed on these clusters. While you create a new Atlas replication policy, do not select **Disable Atlas metadata replication** check-box.

## Snapshot replication between HDP clusters

You can optionally enable HDFS snapshots for replication in Data Lifecycle Manager. Understanding how snapshots work, and some of the benefits and costs involved, can help you to decide whether or not to enable snapshot replication.

Understanding HDFS Snapshots

HDFS snapshots are read-only point-in-time copies of the filesystem. You can enable snapshots on the entire filesystem, or on a subtree of the filesystem. For DLM, you enable snapshots at a dataset level.

Enabling snapshots on a folder requires HDFS admin permissions, because it impacts the NameNode. When you enable snapshots, all subdirectories are automatically enabled for snapshots as well. So when you create a snapshot copy of a directory, all content in that directory, including subdirectories, is included as part of the copy. If a directory contains snapshots but the directory is no longer snapshot-enabled, you must delete the snapshots prior to enabling the snapshot capability on the directory.

Snapshots must be taken on the highest-level parent directory that is snapshot-enabled. Snapshot operations are not allowed on a directory if one of its parent directories is already snapshot-enabled (snapshottable) or if descendants already contain snapshots.

For example, in the directory tree image below, if directory-1 is snapshot-enabled but you want to replicate subdirectory-2, you cannot select only subdirectory-2 for replication. You must select directory-1 for your replication policy.

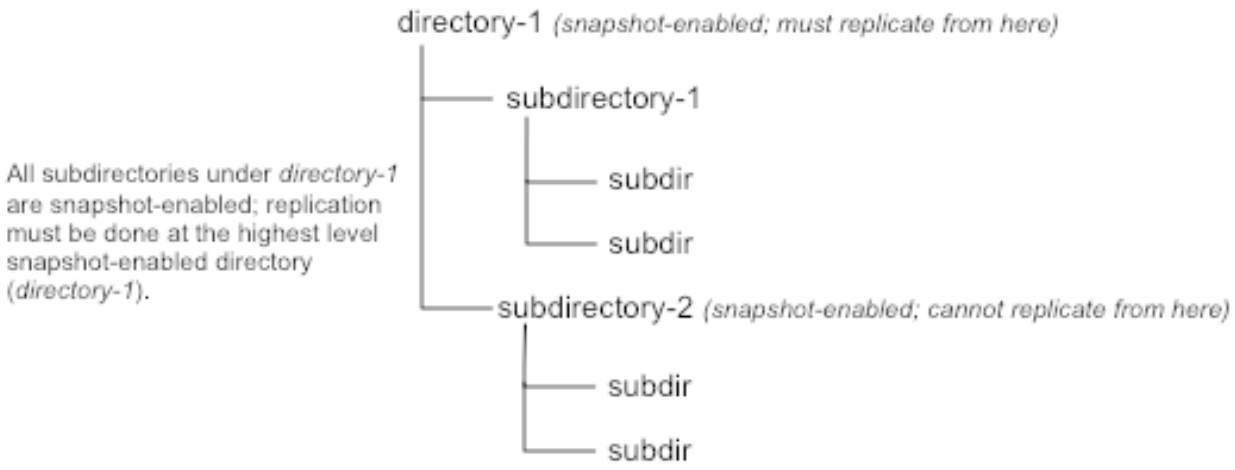


There is no limit to the number of snapshot-enabled directories you can have. A snapshot-enabled directory can accommodate 65,536 simultaneous snapshots.

Blocks in datanodes are not copied during snapshot replication. The snapshot files record the block list and the file size. There is no data copying.

When snapshots are initially created, a directory named .snapshot is created on the source and destination clusters, under the directory being copied. All snapshots are retained within .snapshot directories. By default, the last three snapshots of a file or directory are retained. Snapshots older than the last three are automatically deleted.

Benefits of snapshots

Snapshot-based replication helps you to avoid unnecessarily copying renamed files and directories. If a large directory is renamed on the source side, a regular DistCp update operation sees the renamed directory as a new one and copies the entire directory.

Generating copy lists during incremental synchronization is more efficient with snapshots than using a regular DistCp update, which can take a long time to scan the whole directory and detect identical files. And because snapshots are read-only point-in-time copies between the source and destination, modification of source files during replication is not an issue, as it can be using other replication methods.

A snapshot cannot be modified. This protects the data against accidental or intentional modification, which is helpful in governance and in meeting disaster recovery (DR) requirements.

Considerations for using snapshots

There is a memory cost to enabling and maintaining snapshots. Tracking the modifications that are made relative to a snapshot increases the memory footprint on the NameNode and can therefore stress NameNode memory.

Because of the additional memory requirements, snapshot replication is recommended for situations in which it is most useful. Such circumstance might include: if you expect to do a lot of directory renaming, if the directory tree is very large, or if you expect changes to be made to source files while replication jobs execute.

**Considerations for HDP 2.6.5**: HDP 2.6.5 release does not support snapshots in the true sense. In case of open files, the snapshots do not have point in time copy of the same and this can lead to data loss. The snapshot-diff based replication is disabled in HDP 2.6.5. If you enable snapshot-based replication in HDP 2.6.5, the snapshots are used as the source, without using the diff-based replication. While using snapshot as a source during replication and if any modification is performed in the source, say file deletion, the replication process shall continue without throwing any file not found exception.

Requirements for snapshot-based replication

You must have HDFS superuser privilege to enable or disable snapshot operations.

Replication using snapshots requires that the target filesystem data being replicated is identical to the source data for a given snapshot. There must not be any modification to the data on the target. Otherwise, the integrity of the snapshot cannot be guaranteed on the target and replication can fail in various ways.

# Replication policy operations

This page provides information about various tasks while running the data replication policy.

## Monitoring replication

Ensure that the frequency is set so that a job finishes before the next job starts. Jobs based on the same policy cannot overlap.

If a job is not completed before another job starts, the second job does not execute and is given the status Skipped. If a job is consistently skipped, you might need to modify the frequency of the job.

## Policies page

You can check job status from several places in the DLM UI.

### Before you begin

### About this task

You can view the status and other information about policies and associated jobs from the Policies page. All jobs (policy instances) can be viewed from this page, regardless of status.

The Policies Page can display up to 200 policies.

### Procedure

1. In the navigation pane, click **Policies**.
2. Click the
   
   
   or
   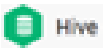
   
   icon to display the type of policies you want to view.
3. Locate the policy associated with the job that you want to view by doing one of the following:
   - Browse the list to find the name of the policy.
   - Enter full or partial terms in the search field.
4. For the policy you located, click
   
   in the Prev Jobs column to open or close the list of jobs associated with the policy.

   A maximum of 10 jobs displays per page.
5. Click
   
   to see the next or the previous list of jobs.

## Overview page

The Overview page displays jobs that are either in progress or have not succeeded. While jobs are executing, they display in the list with a status of In Progress. If the job succeeds, it disappears from the list. Successful jobs can be viewed from the Policies page.

### Procedure

1. In the navigation pane, click **Overview**.
2. Browse the Issues & Updates list to locate the policy for the job you want status for.
3. View the Job Status column for the policy.
4. If the job did not succeed, click
   
   next to the job status to view the job log.
5. Optionally, see information about previous job runs:
   a) Click the dots in the Policy History column.
      The policy displays in the Policies page.
   b) Click the dots in the Prev Job column.
      A list of jobs related to the selected policy displays, showing up to the last 10 jobs.

## Notifications page

You can view the ongoing and completed activities on the notification page in the DLM UI.

### Before you begin

You must use the DLM Infrastructure Admin role to perform this task.

### Procedure

1. From any page in Data Lifecycle Manager, click
   
   to display the last five job alerts.
2. From the Notifications dialog box, click **View All** to open the Notifications page, showing all previous notifications.

## Viewing replication logs

After creating the replication policy, you can view the running logs of the executed replication job.

### About this task

You must use the DLM Infrastructure Admin role to perform this task.

### Procedure

1. In the navigation pane, click Policies.
2. Choose the type of policy you want to view. For example, HDFS or HIVE.
3. Select the policy associated with the replication job and click the
   
   icon.
4. Select View Log.
   The association replication job log window is displayed.
5. Optional: Download or copy or open the log file.

## Tuning replication policy (advanced options)

Specify bandwidth per map, in MBps. Each map 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 bandwidth used tends towards the specified value.

Queue Name

If you are using Capacity Scheduler queues to limit resource consumption, enter the name of the YARN queue for the cluster to which the replication job will be submitted.

Maximum Bandwidth

You can adjust this setting so that each map task is throttled to consume only the specified bandwidth so that the net bandwidth used tends towards the specified value. The default value for the bandwidth is 100 MB per second.

Maximum Maps

Use this option to set the maximum number of map tasks (simultaneous copies) per replication job.

The Advanced Settings attributes are applied only during DLM replication jobs that are based on DistCp functionality.

## Suspend data replication

When you create and run the replication policy, during the course of the replication, you can suspend data replication.

### About this task

You must use the DLM Infrastructure Admin role to perform this task.

### Procedure

1. In the navigation pane, click Policies.
2. Choose the type of policy you want to temporarily stop. For example, HDFS or HIVE.
3. Select the policy associated with the replication job and click the

   

   icon.
4. Click Suspend.
   A message appears if the selected replication policy is to be suspended.
5. Select Yes.
   A confirmation message notifies that the selected policy is suspended.

## Activate data replication

You can activate an already suspended replication policy.

### About this task

You must use the DLM Infrastructure Admin role to perform this task.

### Procedure

1. In the navigation pane, click Policies.
2. Choose the type of policy you want to activate. For example, HDFS or HIVE.

3. Select the policy associated with the replication job and click the

   

   icon.
4. Click Activate.
   A message appears if the selected replication policy is to be activated.
5. Select Yes.
   A confirmation message notifies that the selected policy is now activated.

## Edit replication policy

You can edit some settings in your 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.

### About this task

The Edit Replication Policy page is not available prior to DLM version 1.1.1.

- You can edit an existing policy, with the following restrictions:
  - Only non-expired policies in active or suspended state can be edited.
  - The start time cannot be modified if the policy has already started.
  - You cannot modify the policy name or the source or destination cluster.
- DLM does not support update of any cluster endpoints (HDFS, Hive, Ranger, or DLM Engine). If an endpoint must be modified, contact Hortonworks support for assistance.

### Before you begin

You must use the DLM Infrastructure Admin role to perform this task.

### Procedure

1. In the DLM navigation pane, click **Policies**.
   The Replication Policies page displays a list of any existing policies.
2. Locate the policy you want to edit and click

   

   (Actions).

   
   
3. Select **Edit** and then modify and save the policy.
   The following options are available to edit:
   - Frequency
   - Start Date (if the policy has not yet run an initial job instance)
   - End Date
   - Start Time
   - Queue Name
   - Maximum Bandwidth
   - Maximum Maps

   If the **Edit** option does not display, verify that the policy status is active or suspended. Expired policies cannot be edited.
4. To edit policy description and key selection, on the Policies Page, click the policy name.

The Policy Settings display.

Key selection is only available for policies that are replicating TDE-enabled data.

5. Click the **Edit** icon next to Description or Source & Destination Encrypted With.

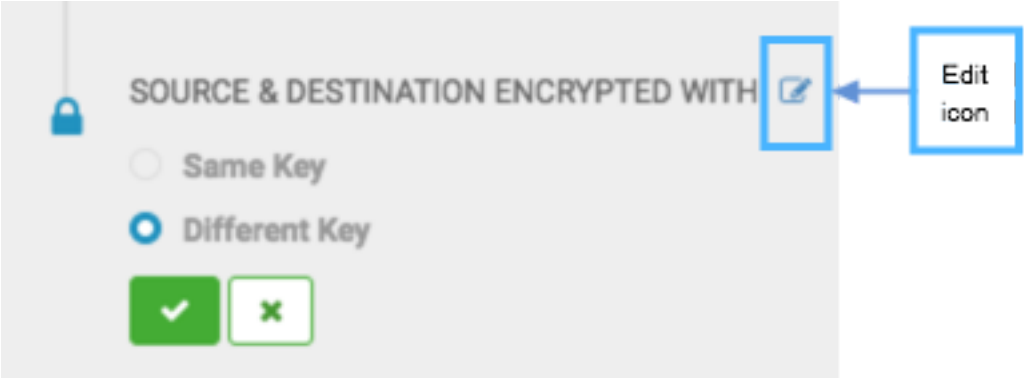


Clicking the Edit icon next to other items in Policy Settings opens the Edit Replication Policy wizard.

6. Click the checkmark to save the change and close the edit option.

### What to do next

View job status to verify that the replication job is running as intended.

## Browsing data directory

Any user with access to the DLM UI has the ability to browse, within the DLM UI, the folder structure of any clusters enabled for DLM.

Therefore, the DPS Admins and the Infra Admins can see folders, files, and databases in the DLM UI that they might not have access to in HDFS. The DataPlane Admin and Infra Admin cannot view from the DLM UI the content of files on the source or destination clusters. Nor do these administrators have the ability to modify or delete folders or files that are viewable from the DLM UI.

# Cloud credentials operations

## Update cloud credentials

You can update cloud credentials based on various factors

- Changes made to a bucket configuration (secret/access keys, bucket name/endpoint, and encryption type) can affect DLM replication policy execution and might require an update to DLM cloud credentials.
- Credential changes are picked up by the next run of the policy. Any policies being run when the credential changes are made could fail, but succeeding runs will pick up the changes.

## Delete credentials

You can delete unwanted credentials from DLP app

- Users can delete cloud credentials, but this triggers failures of any policies based on the deleted cloud credentials.

- You must delete the DLM cloud policies associated with the deleted credentials and recreate the policies with the new credentials. You can view a list of policies associated with specific credentials on the **Cloud Credentials** page.

## Unregistered credentials

Unregistered credentials can impact replication process

- Unregistered credentials in DLM are credentials associated with a cluster node that do not have updated credentials.
- An example of how this can arise is, if a node was down when the credentials were changed on a bucket, and when the node is brought up it still has the old credentials.

# Miscellaneous

You must update the cluster endpoint.

## Update Cluster Endpoint

A DLM endpoint server is present for each cluster on the DataPlane that has DLM Engine installed. As an administrator, you can change the specific configurations on Ambari to update any cluster endpoint and ensures that it works with DLM.

### Before you begin

You must use the DLM Infrastructure Admin role to perform this task.

### Procedure

1. Log in to DataPlane services.
2. On the navigation pane, click **DATA LIFECYCLE MANAGER**.
3. Click **Clusters**.

   You can view the DLM Engine clusters on the **Clusters** page.
4. Click the **Sync** button to synchronize the changes between the Ambari cluster and the DLM Engine.

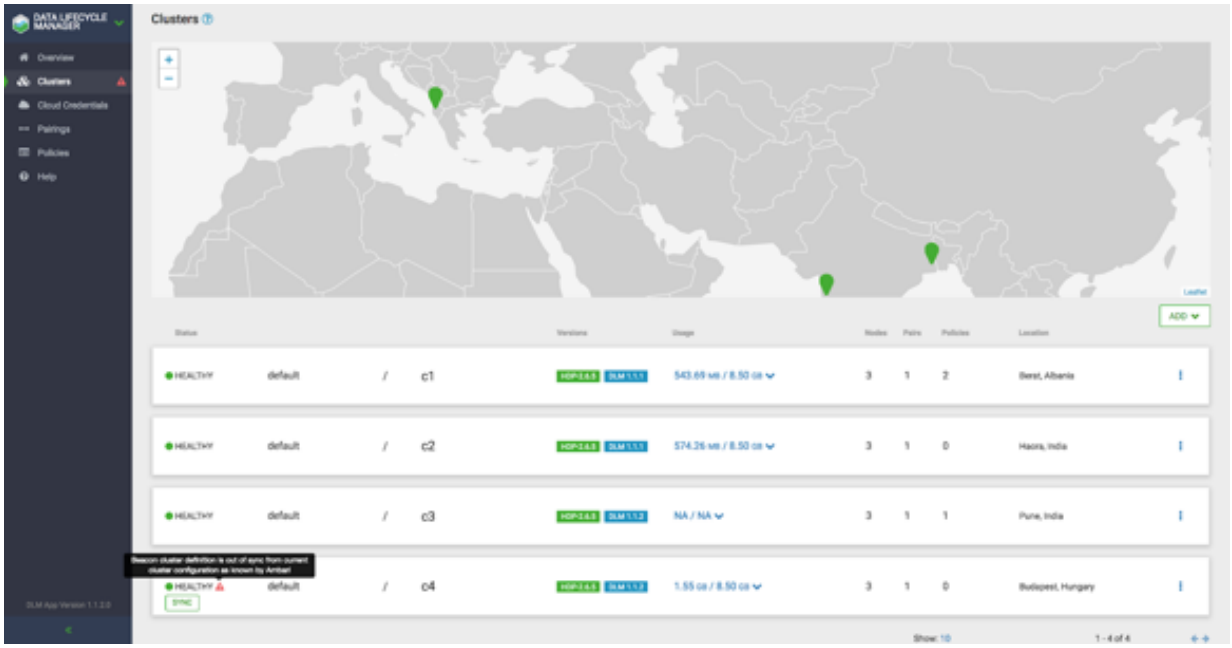


# Failing Over Manually

If a source cluster used in a replication policy is offline and will not be brought online for an extended period, you should manually fail over the destination cluster to serve as the new source. After failover, the new source cluster will receive read and write requests. You might also want to designate a new destination cluster to which data will be copied from the new source.

## Make the destination cluster the new source

If the source cluster becomes unavailable for an extended period, you can configure the destination cluster to serve as the new source. Read and write requests from clients will then be redirected from the old source to the new source cluster.

### Before you begin

You must be logged in as Infra Admin to perform this task.

You need the name of the cluster that is offline.

### Procedure

1. Log in to the DPS UI as Infra Admin.
2. Access the DLM UI by clicking the DPS icon in the upper left of the page and then clicking the Data Lifecycle Manager icon.
3. Identify the set of replication policies for which the offline cluster is the source in a replication relationship.
   a) Click **Policies** in the navigation pane.

b) In the **Filter** field, type the name of the offline cluster.

A list appears that displays the cluster name as a source or a destination cluster.

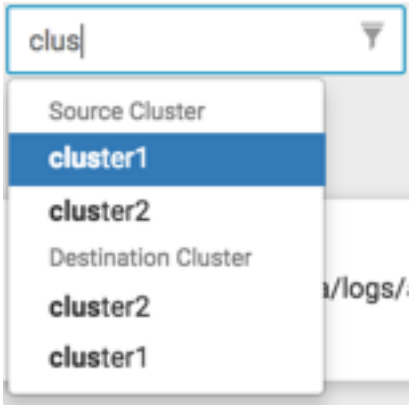


c) From the list, select the cluster name under Source Cluster.

The page content shows only the policies that use the selected cluster as the source for replication.

4. Delete all replication policies that use the offline cluster as the replication source.

a) At the end of each row in the policies list, click the

(Actions) icon.

b) Click **Delete** in the drop-down menu, and then click **OK** to confirm deletion.

If a replication policy is in the process of running a job, the job aborts when you delete the policy.

**Important:** After a replication policy is deleted, it cannot be retrieved.

### What to do next

If the Ranger deny policy is enabled, remove the deny policy that is on the destination cluster.

## Remove the Ranger deny policy

If the Ranger deny policy is enabled, you must remove the deny policy that is on the destination cluster so that DLM can access the target data to be retrieved.

### Before you begin

You must be logged in as Ambari Admin to perform this task.

### Procedure

1. Determine if the Ranger deny policy is enabled.
   a) Navigate to the Ambari UI.
   b) In the services list, click **DLM Engine**.
   c) Click **Configs**>**Advanced**.
   d) Scroll to the parameter beacon.ranger.plugin.create.denypolicy and verify if the **Ranger Deny Policy** is enabled or disabled.
2. If the **Ranger Deny Policy** is enabled, you must disable it.
   a) Log in to the destination cluster, access Ranger, and then navigate to Ranger admin resource policies.
   b) Identify Ranger policies that start with “<sourcecluster>_beacon deny policy for” and remove the deny condition on the policies.

## Activate a new destination cluster

If you have not prepared a cluster in advance to serve as an alternate destination in a failover scenario, then you must install the DLM Engine, configure the clusters for use by DLM, and pair the clusters before you can create new replication policies and begin copying data to the new destination.

### Before you begin

You must have the name of the cluster you want to configure as the new destination.

### Procedure

1. Identify the Ambari-managed cluster to use as the new destination.
2. Install the DLM Engine on the new destination, if it is not already installed. For more information, see Installing DPS Services, Engines, and Agents.
3. Follow the instructions in Setting Up the DPS Services for the following tasks, as needed:
   - *Register Clusters with DPS*
   - *Enable Services*
4. Pair the clusters you are using as source and destination, if they are not already paired. For more information, see Cluster pairing.
5. Ensure that the HDFS folders or Hive databases to be copied either do not exist or are empty on the new destination cluster.

   This is required prior to bootstrapping data from the source cluster to the destination cluster. Otherwise, the initial copy job fails.
6. Create and submit new replication policies between the source and destination clusters.

   The first time a new policy is submitted, the entire contents of the source dataset is copied to the destination. Depending on the size of each dataset, these initial bootstrap copies can take a significant amount of time. After the initial copy, subsequent copies are incremental.

# DLM version Information

As a **Infra Administrator** or **DLM Administrator**, you can view various version-related details on the DLM user interface.

You can view the following details on the DLM UI:

- DLM Engine version
- HDP version on each cluster

**Note:** You can view the DLM Engine and HDP for each cluster on the following pages: **Clusters**, **List Pairings**, and **Create Pairings**.

# Tuning DLM Engine

You can tune the DLM Engine for tasks such as running multiple concurrent policies and handling multiple files.

### Run Multiple Concurrent Policies

Perform the following steps to run multiple concurrent policies in DLM:

1. Log in to Ambari.

2. Set the beacon_quartz_thread_pool property to a value greater than the number of policies required to run concurrently.

#### Handle Multiple Files

For the DLM Engine to handle multiple files that are listed, ensure that it has sufficient memory.

Under /etc/beacon/conf/beacon_env.ini, set the heap value as applicable for BEACON_SERVER_HEAP paramenter. The default value is Xmx2048m.

The default value is sufficient to handle one million files on source dataset. If you have higher number of files in source dataset, change the heap value accordingly.

#### HDFS replication fails with connection refused error on the HA cluster

In a HA-enabled clusters, two parameters, namely, yarn.resourcemanager.connect.max-wait.ms and yarn.resourcemanager.connect.retry-interval.ms are used to connect to Resource Manager (RM). The default values are -1 and 30 seconds respectively. In case of the first parameter, the client waits indefinitely to connect to the RM. And in case of the other parameter, YARN waits for upto 30 seconds to connect with each RM.

To overcome the connectivity problems, in Beacon, a new parameter called yarn_rm_connect_timeout is added, which can overwrite yarn.resourcemanager.connect.max-wait.ms value. The default value of yarn_rm_connect_timeout is set to120 seconds in Beacon. It ensures that all four RMs are tried. You can tune this parameter based on your setup.

# Troubleshooting DLM

To verify that your environment meets the requirements for DataPlane platform, see the DP Support Matrices.

## Ranger UI does not display deny policy items

Deny policy not getting displayed.

If you need to view deny policy details related to a DLM replication policy, you need to use the Ranger UI. However, when a policy with deny conditions is created on Ranger-admin in a replication relationship, the Policy Details page in Ranger does not display the deny policy items. To make the policy visible, update the respective service-def with enableDenyAndExceptionsInPolicies="true" option.

Refer to section "2.2 Enhanced Policy model" in https://cwiki.apache.org/confluence/display/RANGER/Deny-conditions+and+excludes+in+Ranger+policies.

## Replication fails with TDE and non-TDE data

Issue with TDE and non-TDE entities.

HDFS Replication fails when some files are encrypted and some are unencrypted. If the source directory is unencrypted, but contains both encrypted and unencrypted subfolders, then replication jobs fail with checksum mismatch error.

Ensure that all folders in a source *root* directory have the same encryption setting (enabled/not enabled or same key).

## Hive data cannot be replicated

If an initial Hive replication (bootstrap) fails in DLM, review the following possible causes and resolutions to try resolving the issue. .

Notification events are missing in the meta store

REPL_EVENTS_MISSING_IN_METASTORE (20016)

Use the drop command to delete the target database and then resume the policy from the DLM App UI.

Target database is bootstrapped from some other path.

REPL_BOOTSTRAP_LOAD_PATH_NOT_VALID (20017)

If you create a new policy and view this error on the first policy run, use the drop command to delete the target database and later resume the policy from the DLM App UI.

Alternatively, when you upgrade to DLM-1.4.0.0 and edit an existing Hive policy to provide the External Table Base Directory path, you might hit with an error on the next policy instance run. In this case, you must delete all the external tables and resume the policy.

Please work with Hortonworks Support team for assistance.

File is missing from both the source and CM path.

REPL_FILE_MISSING_FROM_SRC_AND_CM_PATH (20018)

Review the DLM Engine logs to locate the REPL DUMP directory, remove the directory, delete (drop) the target database, and then resume the policy from the DLM App UI.

Either the dump directory does not exist or it is not accessible

REPL_LOAD_PATH_NOT_FOUND (20019)

If the dump location does not exist, you can resume the policy and the DLM Engine creates a new dump.

If the directory is not accessible, you need to set the required permissions.

The source for the replication (repl.source.for) is not set in the database properties.

REPL_DATABASE_IS_NOT_SOURCE_OF_REPLICATION (20020)

On the source database, use DESC DATABASE EXTENDED <db_name> to determine if the parameter repl.source.for is set with the policy name.

If the policy is scheduled and the above parameter is not set, then set the parameter using ALTER DATABASE <db_name> SET DBPROPERTIES ('repl.source.for'='<policy_name>').

Then resume the policy from the DLM App UI.

## Hive policy suspension

When you have created a Hive replication policy, the policy can get suspended.

Some of the possible scenarios where your Hive replication policy can get suspended:

- If you upgrade the source cluster from HDP 2.6.5 to HDP 3.x cluster, the already running policy can get suspended.
- If you upgrade the target cluster to HDP 3.x, and later convert the same cluster to any cloud-based cluster, the running policy can get suspended.

## Instance of a policy stuck in a running state

Policy stuck in a running state.

If your policy is stuck in a running state because of some unknown exceptions, you must restart the DLM engine using Ambari. This process would in turn handle the failure scenarios.

**Note:** If a database failure is detected, you must first get the database service up and running.

## Hive replication failure

Hive replication does not succeed.

Hive replication fails with an error message: This operation is not allowed on source cluster: <ClusterOne>. Try it on target cluster: <ClusterTwo>

If the Hive warehouse directory on target cluster is changed from HDFS to Cloud storage, you must Sync the cluster in DLM UI. DLM UI must be aware about the cluster changes.

## About requested events missing in Notification Log table

You must consider all these factors before setting the value of metastore.event.db.listener.timetolive parameter.

The metastore.event.db.listener.timetolive configuration parameter is used to control the time for which an event will be kept in the database listener queue or the backing RDBMS. Note that, if the configuration value is set too high, the number of events in the queue will increase and can impact performance in terms of normal operation. If the value is set too low, the events might get deleted from the queue, before replication could read it and thus could cause incremental replication to fail. In such a scenario, you must bootstrap the system again to get back to the consistent replicated state.

The value of metastore.event.db.listener.timetolive parameter must be large enough to avoid cleaning up of events, when replication of previous events is in progress. During the bootstrap phase, events added once bootstrap starts should be present for the next incremental to succeed. As the bootstrap is performed for whole database, it might consume more time. In case of incremental load, if the replication frequency is too low, incremental load gets triggered that will have large number of events to replicate. This may increase the time required to execute the load.

So, while setting the parameter value, the replication trigger frequency should be taken into consideration to avoid events getting cleaned up before replication finishes. The replication time depends on many factors like the amount of data to be replicated, bandwidth between the clusters, and number of objects like partitions, table, and functions present in the database. For replicating to a cloud-based cluster, the time taken is more as the file system operations takes longer in cloud file system than in HDFS.