

Cloudera Data Warehouse on premises Environments

Date published: 2020-08-17

Date modified: 2025-11-08

CLOUDERA

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

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

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

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

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

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

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

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER'S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

About on premises environments.....	4
View environment details.....	5
Schedule pods on dedicated nodes (node tainting).....	5
Cloudera Data Warehouse backup and restore.....	6
What is backed up.....	6
What is restored.....	7
Backup-restore prerequisites.....	8
Backing up Cloudera Data Warehouse.....	8
Monitoring Hue database backup.....	10
Monitoring Cloudera Data Visualization database backup.....	10
Restoring Cloudera Data Warehouse.....	11


About on premises environments

An environment is a logical entity that represents the association of your on premises user account with compute resources. Using the compute resources of the environment, you can provision and manage Cloudera Data Warehouse, Cloudera Data Engineering, or Cloudera AI workloads.

For a Cloudera on premises deployment, you can configure the environment on an OpenShift Container Platform or an Embedded Container Service (ECS) cluster. Deploying Cloudera Data Services on premises on OpenShift requires you to deploy and manage the Kubernetes infrastructure. ECS service creates and manages an embedded Kubernetes infrastructure for use with the Cloudera on premises Experiences. To deploy Cloudera Data Services on premises on ECS clusters, you only need to provide hosts on which you want to install the service and Cloudera Manager sets up an ECS cluster using an embedded Kubernetes service and provides a framework to manage and monitor the cluster.

Much of the installation and deployment configuration for private cloud is performed before you register environments using the Cloudera Management Console. However, in addition to the steps described in the installation guide, you must perform additional configurations to activate an environment so you can use it with the Cloudera Data Warehouse service.

The following diagram shows the components of Cloudera on premises:



For more details about environment requirements and how to register an environment so Cloudera can access Kubernetes clusters on your OpenShift or ECS deployment, see [Cloudera on premises Environments](#).

Related Information

[Activating OpenShift environments](#)

[Activating Embedded Container Service environments](#)

Viewing environment details

You can view information about an activated environment, such as when it was created and last updated, or how many Database Catalogs and Virtual Warehouses use the environment from the Cloudera Data Warehouse web interface.

Procedure

1. Log in to the Cloudera Data Warehouse service as DWAdmin.
2. Go to the Environments tab on the [Overview](#) page.
3. Locate the environment you want to view and then click the  Edit .
Environment details for the selected environment are displayed.

Scheduling executor pods on dedicated worker nodes in Cloudera Data Warehouse on premises

To optimize resource utilization on your cluster and improve query performance, Cloudera Data Warehouse allows you to schedule Hive and Impala executor and coordinator pods on dedicated executor nodes that are tainted for Cloudera Data Warehouse and equipped with more local storage, either using SSDs or NVMEs.

About this task

Nodes that have local storage are prioritized for Cloudera Data Warehouse query executor pods, preventing their use by other Cloudera Data Warehouse or data service pods. By default, the Use dedicated nodes for executors option is disabled. When you enable this option, Hive and Impala executor and coordinator pods can only be scheduled on the labeled worker nodes. Make sure you label these nodes in advance.

This task describes how to enable the Use dedicated executor node option on an existing environment:

Before you begin

On Cloudera Embedded Container Service, you must dedicate the Cloudera Embedded Container Service cluster nodes for specific workloads using Cloudera Manager, as described in the *Dedicating ECS nodes for specific workloads* section.

On OpenShift Container Platform (OCP), you must dedicate the OCP cluster nodes for specific workloads using the `kubectl taint` command, as described in the *Dedicate an SSD node for Cloudera Data Warehouse workloads* section.

Procedure

1. Log in to the Cloudera Data Warehouse service as a DWAdmin.
2. On the left navigation pane, go to [Admin Features Advanced Configuration](#) .
3. In Advanced Configuration, select Use dedicated executor node checkbox.
4. Click the Update button.

Related Information

[Dedicating ECS nodes for specific workloads](#)

[Dedicating OCP nodes for specific workloads](#)

[Dedicate a SSD node for Cloudera Data Warehouse workloads](#)

[Activating OpenShift environments](#)

[Activating Embedded Container Service environments](#)

Backup and restore in Cloudera Data Warehouse on premises

The backup/restore feature saves your environment parameters, making it possible to recreate your environment with the same settings, URL, and connection strings you used in your previous environment. Learn about the different backup and restore methods, objects and configurations that are included in the backup, and the backup method to use in different scenarios.

For cluster maintenance and integrity, you can back up the Cloudera Data Warehouse cluster including its Kubernetes objects, persistent volumes, autoscaling configuration, and so on for various Cloudera Data Warehouse entities such as the environment, Database Catalog, Virtual Warehouses, and Data Visualization instance. You can recreate the deployment by restoring the backed-up configurations and settings, as part of the planned maintenance or in a disaster recovery scenario.

Available backup and restore methods

There are two ways to backup and restore the Cloudera Data Warehouse service:

- [Using Data Recovery Service](#)
- [Using the Cloudera Data Warehouse's Cloudera CLI cluster management commands](#)

Both these methods use CDP CLI commands. The Cloudera Management Console also provides a graphical user interface to perform backup and restore operations. See [Data Recovery Service automatic backups](#) and [Using Data Recovery Service with Cloudera Data Warehouse](#).

Deciding which backup method to use

Use Data Recovery Service to back up all Cloudera Data Services on premises environment configurations for an ECS or Cloudera Data Services on premises upgrade. Data Recovery Service takes a snapshot of the namespace on the Kubernetes cluster. You can also choose Data Recovery Service if some Cloudera Control Plane service fails and you need to reinstall it with a previously preserved configuration.

If you just want to back up your Database Catalog or Virtual Warehouse configurations to recreate Cloudera Data Warehouse from an earlier configuration, use Cloudera Data Warehouse's backup and restore commands.

Understanding what is backed up

There are differences in the cluster configurations and objects that get backed up when you create backups using Data Recovery Service or Cloudera Data Warehouse's cluster management Cloudera CLI commands.

Use Data Recovery Service to create on-demand backups of the Data Warehouse namespace, including Kubernetes objects, persistent volumes, and so on. See [Using Data Recovery Service with Cloudera Data Warehouse](#). The output is backupCRN; where CRN is Customer Resource Number, a Cloudera-specific identifier for an environment.

Use the Cloudera Data Warehouse's Cloudera CLI cluster management commands to create a backup from the configuration and settings, including all the connected Database Catalog, Virtual Warehouses, and Cloudera Data Visualization instances. The following settings are included when you back up using the Cloudera Data Warehouse's Cloudera CLI cluster management commands:

- All environment activation settings:
 - Storage Class (Openshift Container Platform)
 - Security Context Constraint (SCC) (Openshift Container Platform)
 - Low Resource Mode
 - Quota Management
 - Resource Pool
 - Dedicated Executor Nodes
 - mTLS client credentials
- All Database Catalog configurations that are associated with the environment. The Hue database is also backed up if you have created one.
- All configurations associated with the Hive and Impala Virtual Warehouses.
- All configurations associated with the Cloudera Data Visualization instances. The database that is associated with the Cloudera Data Visualization instance is also backed up so that the visualizations, dashboards, data connections, and so on are preserved.

Related Information

[Using DRS with Cloudera Data Warehouse](#)

Understanding what is restored

The restore operation restores the state of the Cloudera Data Warehouse service depending on the backup method.

If you have backed up Cloudera Data Warehouse using the Data Recovery Service Cloudera CLI command `dw create-backup`, then the `dw restore-backup` command restores the state of the Cloudera Data Warehouse service from the backup.

If you have backed up Cloudera Data Warehouse using Cloudera Data Warehouse's Cloudera CLI cluster management commands (`dw backup-cluster`), then based on the existing state of the cluster, the restore process creates a workflow plan that decides whether to create, or skip the restore of a component. This workflow plan is returned as the response of the restore command, so you can see which components will be created, updated, or skipped during the restore.

About the restore command

Cloudera Data Warehouse's restore command is as follows:

```
cdp dw restore-cluster --cluster-id <value> --data <value> [--cli-input-json <value>] [--generate-cli-skeleton]
```

The “data” field in the output of the `dw backup-cluster` command contains a base64-encoded zip file containing the backup data of the cluster which includes the environment activation settings, Virtual Warehouse and Cloudera Data Visualization settings and configuration, and locations of Cloudera Data Visualization and Hue database backups on HDFS.

You can use the Cloudera Data Warehouse's `dw restore-cluster` command in one of the following ways:

- By passing the environment's Cloudera resource name (crn) to activate the cluster from the backup file and restore all the entities and database contents.
- By passing an activated environment identifier to restore all the entities and database contents to the running environment. This method is useful when you need to change activation parameters, but it requires manual reactivation.

When you run the `dw restore-cluster` command, Cloudera Data Warehouse:

1. Activates the environment using the settings from the backup and waits for the infrastructure to be created
2. Creates a default Database Catalog
3. Updates the Database Catalog configuration to apply custom configurations

4. Starts the Hue database restore job in the database catalog namespace asynchronously
5. Deploys the Virtual Warehouse instances
6. Deploys the Cloudera Data Visualization instances
7. Starts the Cloudera Data Visualization restore job in the individual namespaces asynchronously. It also restores the database associated with the Cloudera Data Visualization instance.



Note: The Cloudera Data Visualization database restore fails for the first time as it tries to drop tables that do not exist. The Cloudera Data Visualization database restore succeeds the second time.

The restore process is designed to be an idempotent process. You can run it multiple times, if needed. If the environment is activated and healthy, you can run the restore operation multiple times to restore the Virtual Warehouse and Cloudera Data Visualization objects. For every restore operation, the Hue database restore will run. This operation overwrites the Hue database contents. If a Virtual Warehouse or a Cloudera Data Visualization object is not present on the cluster, but the backup file contains it, it is restored to the cluster. If such an entity is already deployed, no changes or configuration updates occur.

Prerequisites for Cloudera Data Warehouse backup and restore in on premises

Learn about the prerequisites that are mandatory for a successful backup and restore event of Cloudera Data Warehouse Data Service.

- You must have an active Cloudera Data Warehouse cluster with one Database Catalog and one or more Hive or Impala Virtual Warehouses.
- Only the default Database Catalog is backed up. Ensure that you do not have custom (non-default) Database Catalogs.
- You have installed and configured CDP CLI version 0.9.99 or higher on a host in your cluster from which you can access Cloudera Data Services on premises.
- You have Cluster Administrator privileges and can access the Cloudera Data Warehouse web interface.
- You must use the same Cloudera Data Warehouse version to restore the files you used during the backup process.

Backing up Cloudera Data Warehouse using the backup-cluster command

Use the backup-cluster command to back up the configuration and settings of all the Database Catalog, Virtual Warehouses, and Cloudera Data Visualization instances within your Cloudera Data Warehouse environment.

Procedure

1. SSH into a host on your cluster from which you can access the Cloudera Data Services on premises cluster.

2. Run the following command to back up the cluster:

```
cdp dw backup-cluster --cluster-id [***CDW-CLUSTER-ID***] [--cli-input-json <value>] [--generate-cli-skeleton]
```

Replace *[***CLOUDERA DATA WAREHOUSE-CLUSTER-ID***]* with the actual cluster ID of your environment. The cluster ID is a unique Cloudera Data Warehouse environment identifier.

`--cli-input-json <value>` and `--generate-cli-skeleton` are optional parameters.

To specify the `--cli-input-json` parameter, you must obtain the skeleton of the JSON file by running the following command:

```
cdp dw backup-cluster --generate-cli-skeleton
```

The output of this command is a JSON object as follows:

```
{  
  "clusterId": ""  
}
```

You can now use the JSON string as a parameter for the `--cli-input-json` command option as follows:

```
cdp dw backup-cluster --cli-input-json '{"clusterId": "[***CDW-CLUSTER-ID***]"}'
```

The output contains the following information:

- `clusterId`: The ID of the cluster, a unique identifier of the Cloudera Data Warehouse environment.
- `operationId`: The ID of the backup operation. You can use the operation ID to query the backup execution details using the CLI.
- `timestamp`: The date of the creation.
- `data`: The backup data and configuration.
- `md5`: The md5 hash of the encoded data. In case the data and its hash are lost, the cluster objects cannot be restored automatically.

3. Save the output in a file.

You need this information during the restoration process.

Results

The Hue backup is stored in the following location:

```
hdfs://cdw-backups/[***TIMESTAMP***]_[***JOB-ID***]/[***ENVIRONMENT-NAME***]/hue-backup
```

The Cloudera Data Visualization backup is stored in the following location:

```
hdfs://cdw-backups/[***TIMESTAMP***]_[***JOB-ID***]/[***DATAVIZ-INSTANCE-NAME***]/viz-backup
```

What to do next

Monitor the database backup jobs. The backup process automatically starts the Hue and Cloudera Data Visualization database backup jobs that you can monitor. Make sure that the database backup jobs complete before destroying the cluster. If you delete the cluster before the jobs are completed, you cannot recover the application contents.

Related Information

[CDP CLI documentation: backup-cluster](#)

Monitoring Hue database backup

To monitor the Hue database backup, log into the cluster and monitor the job status under the Database Catalog namespace.

About this task

The backup-cluster command starts a job to load the database dump file, but does not wait for the job to complete. If you have a large database, the job can take up to an hour to complete. Ensure you allow enough time for the job to succeed.



Note: The job that backs up the Hue database fails on OpenShift Container Platform. This is a known issue in Cloudera Data Services on premises 1.5.4. You must manually backup the Hue database.

Procedure

1. SSH in to a cluster host as an Administrator.
2. Run the following command to fetch the details of the backup job:

```
kubectl get jobs -n [***DATABASE-CATALOG-ID***]
```

Following is a sample output:

```
$ kubectl get jobs -n warehouse-1692037411-96hk
  NAME          AGE      COMPLETIONS   DURATION
  hue-backup-ede2b8bd-1d53-4d23-a0f9-87d8ec658f74   1/1      11s
  hue-query-processor-db-create-job                   1/1      8s
  113s
  42h
```

Results

The Hue backup is stored in the following location:

```
hdfs://cdw-backups/[***TIMESTAMP***]_[***JOB-ID***]/[***ENVIRONMENT-NAME***]/hue-backup
```

Monitoring Cloudera Data Visualization database backup

To monitor Cloudera Data Visualization database backup, log into the cluster and see the job status under the Cloudera Data Visualization namespace.

About this task

The backup-cluster command starts a job to create the database dump file, but it does not wait for it to complete. In case your database size is large, it can take up to 20 minutes for the job to complete. Make sure you leave enough time for the job to succeed.

Procedure

1. SSH in to a cluster host as an Administrator.
2. Run the following command to fetch the details of the backup job:

```
kubectl get jobs -n [***DATA-VISUALIZATION-ID***]
```

Following is a sample output:

```
$ kubectl get jobs -n viz-1692216942-fc2g
```

NAME	COMPLETIONS	DURATION
AGE viz-backup-d874515a-be7e-4902-ac75-269c14f9580c 10m	1/1	3m3s
viz-webapp-vizdb-create-job 99m	1/1	57s

Results

The Cloudera Data Visualization backup is stored in the following location:

```
hdfs://cdw-backups/[***TIMESTAMP***]_[***JOB-ID***]/[***DATAVIZ-INSTANCE-NAME]/viz-backup
```

Restoring Cloudera Data Warehouse using the `restore-cluster` command

You can reactivate the entire Cloudera Data Warehouse environment, which includes your cluster with all settings of the environment that you backed up, using the CLI.

Before you begin

Ensure that you have completed the backup process and deactivated the environment in Cloudera Data Warehouse to ensure successful cluster restoration. During downtime, you must prevent end-users from accessing the cluster.



Caution: You lose any manual modification of the Kubernetes objects or configurations when you bring down the cluster. For example, third-party integrations with Protegrity. Modifications applied using the Cloudera Data Warehouse web interface and settings defined during creation are preserved.

Perform the following steps to deactivate the environment in Cloudera Data Warehouse:

1. Log in to the Cloudera Data Warehouse service as a DWAdmin.
2. Go to the Environments tab on the [Overview](#) page, and click Deactivate corresponding to the environment you want to deactivate.

Procedure

1. SSH into a host on your cluster from which you can access the Cloudera Data Services on premises cluster.
2. Run the following command to restore the cluster:

```
cdp dw restore-cluster --cluster-id [***CLUSTER-ID***] --data [***DATA-STRING-FROM-BACKUP***]
```

Replace `[***CLOUDERA DATA WAREHOUSE-CLUSTER-ID***]` with the actual cluster ID of your environment. The cluster ID is a unique Cloudera Data Warehouse environment identifier.



Important: If the environment is not activated, provide the Environment CRN as the value for `[***CLOUDERA DATA WAREHOUSE-CLUSTER-ID***]`. Provide the unique Cluster ID only if the environment is in a Running state.

Replace `[***DATA-STRING-FROM-BACKUP***]` with the actual data string from the output of the backup.

`[--cli-input-json <value>]` and `[--generate-cli-skeleton]` are optional parameters.

To specify the `--cli-input-json` parameter, you must obtain the skeleton of the JSON file by running the following command:

```
cdp dw backup-cluster --generate-cli-skeleton
```

The output of this command is a JSON object as follows:

```
{
```

```
        "clusterId": ""  
    }
```

You can now use the JSON string as a parameter for the `--cli-input-json` command option as follows:

```
cdp dw backup-cluster --cli-input-json '{"clusterId": "[***CDW-CLUSTER-ID***]" }'
```

The output contains the following information:

- `clusterId`: The ID of the cluster, a unique identifier of the Cloudera Data Warehouse environment.
- `operationId`: The ID of the backup operation. You can use the operation ID to query the backup execution details using the CLI.
- `action`: The action to be taken on the cluster. Possible actions are *Create* and *Skip*.
- `message`: The description of the cluster action.
- `dbcRestorePlans`: Information about the restore-plan of the Database Catalogs.
 - `item`: Status of the entity after the restore operation.
 - `ref`: The reference of the entity in the backup data.
 - `id`: The ID of the entity.
 - `action`: The action associated with the plan. Possible actions are *Create*, *Update*, and *Skip*.
 - `message`: The description of the plan.
- `hueRestorePlans`: Information about Hue's restore plan.
 - `item`: Status of the entity after the restore operation.
 - `ref`: The reference of the entity in the backup data.
 - `id`: The ID of the entity.
 - `action`: The action associated with the plan. Possible actions are *Create*, *Update*, and *Skip*.
 - `message`: The description of the plan.
- `hiveRestorePlans`: Information about Hive Virtual Warehouses' restore plan.
 - `item`: Status of the entity after the restore operation.
 - `ref`: The reference of the entity in the backup data.
 - `id`: The ID of the entity.
 - `action`: The action associated with the plan. Possible actions are *Create*, *Update*, and *Skip*.
 - `message`: The description of the plan.
- `impalaRestorePlans`: Information about Impala Virtual Warehouses' restore plan.
 - `item`: Status of the entity after the restore operation.
 - `ref`: The reference of the entity in the backup data.
 - `id`: The ID of the entity.
 - `action`: The action associated with the plan. Possible actions are *Create*, *Update*, and *Skip*.
 - `message`: The description of the plan.
- `vizRestorePlans`: Information about the restore plan of the Cloudera Data Visualization Apps.
 - `item`: Status of the entity after the restore operation.
 - `ref`: The reference of the entity in the backup data.
 - `id`: The ID of the entity.
 - `action`: The action associated with the plan. Possible actions are *Create*, *Update*, and *Skip*.
 - `message`: The description of the plan.

Results

After several minutes, the environment is activated, and the Virtual Warehouses are created in the new cluster and attached to the Database Catalog. The Virtual Warehouse and Cloudera Data Visualization IDs will be changed.



Note: If you have enabled the Use deterministic namespace names option from the **Advanced Configuration** page, then the Virtual Warehouse and Cloudera Data Visualization IDs do not change after the restore operation.

Database Catalog's ID will change regardless of whether you have enabled the Use deterministic namespace names option.

What to do next

Adjust the Cloudera Data Visualization connection settings to point to the new Virtual Warehouse(s) if the IDs have changed.

Related Information

[CDP CLI documentation: restore-cluster](#)