

Managing Classic Clusters

Date published: 2019-08-22

Date modified:



Legal Notice

© Cloudera Inc. 2024. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Managing classic clusters.....	4
Enabling admin and user access to classic clusters.....	4
Prerequisites for adding classic clusters.....	4
Adding classic clusters with CCMv1.....	6
Adding an HDP cluster.....	6
Adding a CDH cluster.....	8
Adding a CDP PvC Base cluster.....	10
Replication Manager use case.....	10
Replication Manager and Data Catalog use case.....	12
Non-transparent proxy.....	14
Troubleshooting classic clusters.....	15
Troubleshooting: Install autossh package.....	17
Adding classic clusters with CCMv2.....	18
Adding an HDP cluster.....	18
Adding a CDH cluster.....	20
Adding a CDP PvC Base cluster.....	22
Replication Manager use case.....	23
Replication Manager and Data Catalog use case.....	24
Non-transparent proxy.....	26
Troubleshooting classic clusters.....	26
Resume the cluster registration.....	28
Delete an unregistered classic cluster.....	28
Managing a classic cluster.....	29
Upgrading a classic cluster from CCMv1 to CCMv2.....	32

Managing classic clusters

You can register CDH, HDP, and CDP Private Cloud Base clusters to CDP and later replicate data and workloads from these clusters to the clusters in your CDP environments. These clusters are called Classic Clusters in CDP.

To add or register CDH, HDP, and CDP Private Cloud Base to CDP:

1. Make sure all the prerequisites are met before adding the classic cluster.
2. Configure the classic cluster to establish connectivity to CDP.
3. On the CDP Management Console, add the classic cluster.

For more information, refer to the following documentation:

Related Information

[Prerequisites for adding classic clusters](#)

[Enabling admin and user access to classic clusters](#)

[Troubleshooting classic cluster registration errors \(CCMv2\)](#)

[Adding classic clusters with CCMv1](#)

[Adding classic clusters with CCMv2](#)

[Resume the cluster registration](#)

[Upgrading a classic cluster from CCMv1 to CCMv2](#)

Enabling admin and user access to classic clusters

In order to grant admin and user access to a classic cluster, you should assign the required roles.

Consider the following when granting access to admins and users of your classic clusters:

- You need to be a ClassicClustersCreator in order to register classic clusters.
- The user who registers a classic cluster gets the Owner role for that classic cluster.
- Once a classic cluster is registered, the following roles can be assigned to a user or a group on that cluster:
- Owner - Grants all permissions on the classic cluster in CDP including the ability to delete them. It does not grant any cluster-level access (such as Cloudera Manager access).
- ClassicClusterAdmin - Grants permission to perform any operation on the cluster, except deleting it. Grants permission to assign access to the cluster to other users.
- ClassicClusterUser - Grants permission to access details of the classic cluster.

The roles are described in detail in Resource roles. The steps for assigning the roles are described in Assigning resource roles to users and Assigning resource roles to groups.

Related Information

[Account roles](#)

[Resource roles](#)

[Assigning resource roles to users](#)

[Assigning resource roles to groups](#)

Prerequisites for adding classic clusters

Make sure you verify the version requirements of your classic clusters and install or configure all of the classic cluster requirements before you try to add or register them to CDP.

Required roles

You need to have the ClassicClustersCreator role to register a classic cluster in CDP.

Check the version compatibility

Make sure that the following version requirements are met.

- For adding HDP clusters:
 - HDP - HDP 2.6.5.50000 plus any required patch version
 - DLM Engine version - 1.7.0.0-x
 - Ambari - 2.7.3 or 2.6.2.2 or 2.6.2
- For adding CDH clusters:
 - CDH - 5.x and 6.x
 - Cloudera Manager - 5.x and 6.x
- For adding CDP Private Cloud Base clusters
 - CDP Private Cloud Base cluster 7.1.1 or later, and its respective Cloudera Manager version; For example, CDP Private Cloud Base cluster 7.1.6 uses Cloudera Manager 7.3.1.



Note: Currently, classic cluster registration only supports CentOS and RHEL operating systems, versions 6, 7 and 8.

Verify the required roles are assigned

- Make sure that the user can log in as the admin user to Ambari in the HDP cluster.
- Make sure the user can log in as the admin user to Cloudera Manager in the CDH cluster or the CDP Private Cloud Base cluster.

Install the required components (HDP only)

Fulfill the following requirements if using HDP. If using CDH or CDP Private Cloud Base, skip this section.

- Make sure that Ambari Metrics is installed in Ambari in the HDP clusters.
- Make sure that Ambari is configured with LDAP.
- Make sure that Knox is installed and the default topology is configured with LDAP.



Note: Knox does not need to be used by other services in the cluster; It is only required for CDP communication.

- Make sure that there is at least one topology in the Knox setup with the same LDAP as Ambari.
- If there are policies restricting access through Knox, make sure that Ranger policies allow communication through Knox.
- Make sure that the cluster is configured with Kerberos.
- Make sure that the user credential used for registering the classic cluster is a valid LDAP user with an admin role in Ambari.

Open ports for CCM

If you would like to use Cluster Connectivity Manager v2 (CCMv2), ensure that outgoing traffic is allowed on the port range 6000-6049 (CCMv1) or port range 443 (CCMv2) on the legacy CDH, HDP, or CDP Private Cloud Base cluster.

After making sure that your clusters meet all the requirements, you can add your CDH, HDP, and CDP Private Cloud Base clusters to CDP.

Related Information

[Cluster Connectivity Manager \(CCM\)](#)

[Using classic clusters with a non-transparent proxy \(CCMv2\)](#)

[Using classic clusters with a non-transparent proxy \(CCMv1\)](#)

Adding classic clusters with CCMv1

If your CDP tenant still uses CCMv1, refer to the following documentation for instructions on how to register your existing on-prem cluster in CDP Public Cloud.



Warning: CCMv1 has been replaced with CCMv2. If you started registering your on-premise cluster with CCMv1, we recommend that you delete it and start a new registration process with CCMv2, as described in [Delete an unregistered classic cluster](#) and [Adding classic clusters with CCMv2](#).

Adding an HDP cluster (CCMv1)

Navigation title: Adding an HDP cluster

You must register clusters with CDP before you can use those clusters with CDP services and components. To ensure optimum security, clusters within the customer environment are not accessible for communication. They have private IP addresses and cannot be accessed outside the firewall. However, to add your cluster to the CDP, a communication line needs to be established.



Warning: CCMv1 has been replaced with CCMv2. If you started registering your on-premise cluster with CCMv1, we recommend that you delete it and start a new registration process with CCMv2, as described in [Delete an unregistered classic cluster](#) and [Adding classic clusters with CCMv2](#).

A reverseSSH tunnel solves the problem by establishing a tunnel from the cluster to CDP. You must download and install AutoSSH and the connectivity install scripts to establish a secure two-way communication channel. The AutoSSH ensures that the connectivity is stable. The connectivity scripts and their installation ensure safe connectivity and communication.

Before you begin



Caution: After you register an HDP cluster in CDP, do not change the cluster name in Ambari. A cluster name change in Ambari does not currently propagate to CDP, which can result in issues when using clusters with CDP clusters and components.

- HDP clusters must be managed by Ambari. Clusters that are not managed by Ambari cannot be registered to CDP.
- HDP clusters must include Knox.
- HDP clusters must include Ranger policy settings
- LDAP/AD must be set up and synced in Ambari

LDAP settings are automatically detected from the default topology setup in Knox. If the default topology does not have the LDAP setup, you will be asked to provide another topology name where you have configured the LDAP. If that topology has LDAP the setup continues. If the LDAP is not configured, you will receive an error message.

- Kerberos must be enabled on the HDP cluster and the LDAP/AD must be set up in the Kerberos authentication so that the same set of LDAP/AD credentials can be used to access Ambari APIs as well as Beacon APIs
- All clusters must meet the requirements identified in [Prerequisites for adding classic clusters](#).

Steps

The process to register an HDP cluster using a reverseSSH tunnel is as follows:

1. Log in to CDP and navigate to the Management Console.
2. Click Classic Clusters in the left navigation panel.
3. Click Add Cluster.

Classic Cluster displays the Cluster Details dialog box.

4. If you are a first time user, under Step 1 in the Register Classic Cluster wizard, click GET STARTED. Classic Cluster then displays the Cluster Details dialog box.

If you are not a first time user, click the ADD CLUSTER button on the right side of the listing page.

5. Click HDP.
6. Provide the following connectivity information for your new cluster, then click CONNECT:

- IP Address
- Port
- Data center



Note: Establishing connection might take more than twenty minutes if you are adding the cluster for the first time.

After Classic Cluster establishes connection, it will highlight Step 2 in the Register Classic Cluster wizard.

7. Start the download and installation process for the SSH connectivity files by clicking the Files button in Step 2 of the wizard.
8. Follow the instructions in the Setup Connectivity Client dialog box: Download and install the ccm-autossh-client, and the ssh_tunnel_setup_files or cluster_connectivity_setup_files onto your new cluster.

If the installation of CCM autossh-client rpm fails with “No package autossh available”, refer to [Troubleshooting: Install autossh package](#).

9. Copy the files to your Knox node or the Knox proxy host in the cluster.
10. Run the install.sh script.

```
./install.sh
```

11. Enter the following information as the install script prompts for it:

- Enter Ambari URL (http(s)://host:[port]):
- Enter Ambari Username:
- Enter Ambari Password:

12. If Knox is not installed on a proxy server, proceed to Step 13.

Classic Cluster sets up the topology for the Knox server and establishes the reverseSSH tunnel.

13. If Knox is installed on a proxy server, Classic Cluster displays the following message:

We discovered that your Knox is installed in HA mode. Please confirm if this node is your proxy node (yes/no):
Enter yes.

Classic Cluster generates XML content that you will need to add to your Knox hosts. Classic Cluster also displays three steps you must perform on all of your Knox hosts:

- a. Copy the generated XML to /usr/hdp/current/knox-server/conf/topologies/cdp_default.xml.

For example:

```
<?xml version = '1.0' encoding = 'utf8'?>
<topology>
  <gateway>
    <provider>
      <role>authentication</role>
      <name>ShiroProvider</name>
      <enabled>true</enabled>
      <param>
        <name>sessionTimeout</name>
        <value>30</value>
      </param>
    </provider>
  </gateway>
</topology>
```

- b. Run `chown knox:knox cdp_default.xml`.
- c. Check the Knox logs/deployment directory to verify that the `cdp_default` topology is deployed.
- d. After you have completed Steps a through c on all of your Knox hosts, press Enter to continue.

This sets up the topology for the Knox server or Knox proxy host and establishes the reverseSSH tunnel.

14. Classic Cluster starts checking the connectivity with the cluster. When the connectivity is successful, proceed to Step 3 in the wizard.

If the connection attempts fail or if there is an error in the connectivity, Classic Cluster displays troubleshooting information in Step 2 of the Registration wizard. Follow the troubleshooting information to fix the connectivity error, then click Test connection.



Note: After you download the files, the ssh files download is disabled. At this point, you can regenerate the ssh files using the regenerate files option. This option comes in handy if you lose the files previously downloaded before you can set the cluster connectivity in the cluster.

15. Click Register in Step 3 of the Registration wizard.
16. In the Cluster Details dialog box, provide the username and password to access the cluster, then click CONNECT.

The user should have Admin access to the customer cluster services.

17. Finish registering the cluster by providing the following information.

- Cluster Location
- Data Center
- Tags (optional)
- Description (optional)

If LDAP is not set up on the default topology, the system will ask for the following additional information:

Enter knox topology name that contains LDAP setup:

18. Click ADD.

Adding a CDH cluster (CCMv1)

Navigation title: Adding a CDH cluster

You must register clusters with CDP before you can use those clusters with CDP services and components. You can register CDH clusters that can be reached only through private IP address using a reverseSSH tunnel.

About this task



Warning: CCMv1 has been replaced with CCMv2. If you started registering your on-premise cluster with CCMv1, we recommend that you delete it and start a new registration process with CCMv2, as described in [Delete an unregistered classic cluster](#) and [Adding classic clusters with CCMv2](#).

To ensure optimum security, clusters within the customer environment are not accessible for communication. They have private IP addresses and cannot be accessed outside the firewall. However, to add your cluster to the CDP, a communication line needs to be established.

A reverseSSH tunnel solves the problem by establishing a tunnel from the cluster to CDP. You must download and install AutoSSH and the connectivity install scripts to establish a secure two-way communication channel. The AutoSSH ensures that the connectivity is stable. The connectivity scripts and their installation ensure safe connectivity and communication.

The high-level steps to register a CDH cluster using a reverseSSH tunnel are as follows:

1. On the CDP Management Console, you enter the private IP address of your cluster and provide the cluster details.
2. You download the AutoSSH rpm from the specified location and the connectivity installation scripts from CDP on to the cluster.
3. You install AutoSSH on the cluster.
4. You register the cluster for performing further operations.

More detailed steps are provided below.

Before you begin



Important: After you register a CDH cluster in CDP, do not change the cluster name in Cloudera Manager. A cluster name change in Cloudera Manager does not currently propagate to CDP, which can result in issues when using clusters with CDP clusters and components.

- CDH clusters must have been created using Cloudera Manager. Clusters that are not managed by Cloudera Manager cannot be registered to CDP.
- All clusters must meet the requirements identified in [Prerequisites for adding classic clusters](#).



Note: The AutoSSH and connectivity install scripts files must be stored in a secure environment.

Steps

Perform the following steps to add a CDH cluster that is not reachable via public IP address.

1. Log in to CDP and navigate to the Management Console.
2. Click Classic Clusters in the left navigation panel.
3. Click Add Cluster.

Classic Cluster displays the Cluster Details dialog box.

4. Click CDH.
5. If your cluster is not reachable by a public network, click My cluster is accessible only in my private network.



Note: Make sure that the Data Center name is different from the Data Center names that have already been registered. If the Data Center name exists, make sure that the combination of the Data Center name and the cluster name is unique. Else, it might result in an error when you try to add a cluster with an existing Data Center-cluster name combination.

6. Provide the connectivity information for your new cluster, then click CONNECT.



Note: Connecting to the new cluster might take more than twenty minutes if you are adding the cluster for the first time.

After Classic Cluster successfully connects to your new cluster, it will highlight Step 2.

7. Start the download and installation process for the SSH connectivity files by clicking the Files button in the Step 2 pane.
8. Follow the instructions in the Setup Connectivity Client dialog box. You need to download the `ccm-autossh-client rpm` file, and the `ssh_tunnel_setup_files` or `cluster_connectivity_setup_files` zip file onto Cloudera Manager host in your new cluster and then:
 - a. Install the AutoSSH rpm (the `ccm-autossh-client` file).
If the installation of CCM `autossh-client rpm` fails with “No package `autossh` available”, refer to [Troubleshooting: Install autossh package](#).
 - b. Unzip the `ssh_tunnel_setup_files` or `cluster_connectivity_setup_files` file. Inside this zip file there is a script `install.sh`
 - c. Run `install.sh` by using `./install.sh` command

Five minutes after you download the ssh files, Classic Cluster starts checking the connectivity with the cluster. When the connectivity is successful, proceed to Step 3 in the UI.

If the connection attempts fail or if there is an error in the connectivity, Classic Cluster displays troubleshooting information in the Step 2 pane. Follow the troubleshooting information to fix the connectivity error, then click Test connection.



Note: After you download the files, the ssh files download is disabled. At this point, you can regenerate the ssh files using the regenerate files option. This option comes in handy if you lose the files previously downloaded before you can set the cluster connectivity in the cluster.

9. Click Register in the Step 3 pane.

10. Provide the username and password of the Cloudera Manager user to access the cluster.

11. Finishing registering the cluster by providing the following information.

- Cluster Location
- Data Center
- Tags (optional)
- Description (optional)

12. Click Submit.

Adding a CDP Private Cloud Base cluster (CCMv1)

Navigation title: Adding a CDP PvC Base cluster

You must add a CDP Private Cloud Base cluster with CDP to register it before you can use those clusters with CDP services and components. You can register CDP Private Cloud Base clusters that can be reached only through a private IP address using a reverseSSH tunnel.



Warning: CCMv1 has been replaced with CCMv2. If you started registering your on-premise cluster with CCMv1, we recommend that you delete it and start a new registration process with CCMv2, as described in [Delete an unregistered classic cluster](#) and [Adding classic clusters with CCMv2](#).

To ensure optimum security, clusters within the customer environment are not accessible for communication. They have private IP addresses and cannot be accessed outside the firewall. However, to add your cluster to CDP, a communication line needs to be established. A reverseSSH tunnel solves the problem by establishing a tunnel from the cluster to CDP. You must download and install AutoSSH and the connectivity-install scripts to establish a secure two-way communication channel. The AutoSSH ensures that the connectivity is stable. The connectivity scripts and their installation ensure safe connectivity and communication.



Note: The AutoSSH and connectivity-install scripts files must be stored in a secure environment.

To register the CDP Private Cloud Base cluster as a classic cluster, you enter the CDP Private Cloud Base cluster details. The Management Console acquires the configuration details from Cluster Connectivity Manager (CCM) service and saves it as ZIP files. You download the ZIP files, install the acquired configurations, and then register the CDP Private Cloud Base cluster as a classic cluster.



Note: After you register a CDP Private Cloud Base cluster in CDP, do not change the cluster name in Cloudera Manager. A cluster name change in Cloudera Manager does not currently propagate to CDP, which can result in issues when using clusters with CDP clusters and components.

All the clusters must meet the requirements identified in [Prerequisites for adding classic clusters](#).

You have two options for registering your CDP Private Cloud Base cluster:

- If you would like to use the CDP Private Cloud Base cluster with Replication Manager, register the cluster using Cloudera Manager.
- If you would like to use the CDP Private Cloud Base cluster with Replication Manager and Data Catalog, register the cluster using Cloudera Manager and Knox.



Important: CDP Private Cloud Base clusters can be used in Data Catalog by registering them using Cloudera Manager and Knox endpoints. Note that this is a technical preview feature and is under development. Do not use this in your production environment. If you have feedback, contact Support by logging a case on the Cloudera Support Portal at <https://my.cloudera.com/support.html>. Technical preview features are not guaranteed troubleshooting and fixes.

For registration steps, see the following documentation:

Adding CDP Private Cloud Base cluster for use in Replication Manager (CCMv1)

Navigation title: Replication Manager use case

After you register the CDP Private Cloud Base cluster as a classic cluster using Cloudera Manager, you can use the classic cluster as a source cluster in Replication Manager.

Before you begin



Warning: CCMv1 has been replaced with CCMv2. If you started registering your on-premise cluster with CCMv1, we recommend that you delete it and start a new registration process with CCMv2, as described in [Delete an unregistered classic cluster](#) and [Adding classic clusters with CCMv2](#).

Procedure

1. Log in to CDP Management Console.
2. Click Classic Clusters.
3. On the Classic Clusters page, click ADD CLUSTER.
4. In the Add Cluster dialog box, go to the CDP Private Cloud Base tab to enter the following details:
 - a) IP address - Enter the IP address of the Cloudera Manager of the CDP Private Cloud Base cluster. The Management Console uses this IP address to identify the cluster for registration purposes.
 - b) Port - Enter the port of the Cloudera Manager of the CDP Private Cloud Base cluster.
 - c) Data center - Enter a unique datacenter name for the CDP Private Cloud Base cluster.
 - d) Select the My cluster runs on HTTPS option if the CDP Private Cloud Base cluster uses HTTPS.
 - e) Clear the Register KNOX endpoint (Optional) option, if selected.
 - f) Click CONNECT.

The Management Console acquires the configuration details from Cluster Connectivity Manager (CCM) service.

5. On the Classic Clusters page, click Files in the Step 2 pane.
6. In the Setup Connectivity Client dialog box, click `ccm-autossh-client`, and `ssh_tunnel_setup_files` or `cluster_connectivity_setup_files` links to download the RPM for AutoSSH from the specified location and the connectivity-installation scripts from CDP, to your machine.
7. In the Setup Connectivity Client dialog box, click Close.
8. In the command line interface, copy the RPM and ZIP files to the host where Cloudera Manager is running.
9. To establish an reverseSSH tunnel between the Cloudera Manager of CDP Private Cloud Base cluster and Management Console, perform the following steps:
 - a) SSH to the Cloudera Manager host.
 - b) Run the following command to install AutoSSH using the `ccm-autossh-client` rpm file:

```
yum --nogpgcheck localinstall [***downloaded ccm_autossh-client-rpm***]
```

If the installation of CCM `autossh-client` rpm fails with “No package `autossh` available”, refer to [Troubleshooting: Install autossh package](#).

- c) Extract the `ssh_tunnel_setup_files.zip` or `cluster_connectivity_setup_files.zip` file to a directory.
- d) Run the install script in the directory using the `./install.sh` command.
- e) Optionally, you can check the CCM service status using `ccm-tunnel status` command to verify whether the SSH tunnel is established.



Note: If you regenerate the script files, you cannot use the previously downloaded `ssh_tunnel_setup_files.zip` or `cluster_connectivity_setup_files.zip` file because the file is no longer valid.

10. On the Classic Clusters page, click Test Connection in the Step 2 pane to verify whether the tunnel connection is successful.
11. Click Register in the Step 3 pane.
12. In the Cluster Details dialog box, enter the Cloudera Manager credentials that has Admin access to Cloudera Manager and the cluster services.
13. Click CONNECT.

14. To complete the registration, enter the following details on the Classic Clusters page:

- Cluster Location - Enter the geographical location of the Data Lake.
- Data Center - Ensure that the data center name is the name that you provided for CDP Private Cloud Base cluster during registration.
- Tags - Optionally, enter the tags for the cluster.
- Description - Optionally, enter a description.

15. Click Add.

Results

You can use the registered classic cluster in Replication Manager.

Adding CDP Private Cloud Base cluster for use in Replication Manager and Data Catalog (CCMv1)

Navigation title: Replication Manager and Data Catalog use case

After you register the CDP Private Cloud Base cluster as a classic cluster using Cloudera Manager and Knox endpoints, you can use the classic cluster in Replication Manager and Data Catalog.

Before you begin



Warning: CCMv1 has been replaced with CCMv2. If you started registering your on-premise cluster with CCMv1, we recommend that you delete it and start a new registration process with CCMv2, as described in [Delete an unregistered classic cluster](#) and [Adding classic clusters with CCMv2](#).

Ensure that the following components and role are available:

- The CDP Private Cloud Base cluster has an active Knox service.
- A proxy to Cloudera Manager through Knox for communication purposes. For more information, see [Proxy Cloudera Manager through Apache Knox](#).
- LDAP is configured in the Cloudera Manager of CDP Private Cloud Base cluster. For more information, see [Configure authentication using an LDAP-compliant identity service](#).
- A minimum of one LDAP user with the Full Administrator role.
- An LDAP-based topology with CM-API, CM-UI, ATLAS, ATLAS-API, RANGERUI, and RANGER services. The topology name is used during the classic cluster registration process.

Disposition: / Status:

Removed link to Knox UI doc that was removed.



Note: If there are policies that restrict access through Knox, then add the topology name to the cdp_default Ranger policy so that the Ranger policies can communicate through Knox.



Important: CDP Private Cloud Base clusters can be used in Data Catalog by registering them using Cloudera Manager and Knox endpoints. Note that this is a technical preview feature and is under development. Do not use this in your production environment. If you have feedback, contact Support by logging a case on the Cloudera Support Portal at <https://my.cloudera.com/support.html>. Technical preview features are not guaranteed troubleshooting and fixes.

Procedure

- Log in to CDP Management Console.
- Click Classic Clusters.
- On the Classic Clusters page, click ADD CLUSTER.

4. In the Add Cluster dialog box, go to the CDP Private Cloud Base tab.
 - a) IP address - Enter the IP address of the Cloudera Manager of the CDP Private Cloud Base cluster. The Management Console uses this IP address to identify the cluster for registration purposes.
 - b) Port - Enter the port of the Cloudera Manager of the CDP Private Cloud Base cluster.
 - c) Data center - Enter a unique datacenter name for the CDP Private Cloud Base cluster.
 - d) Select the My cluster runs on HTTPS option if the CDP Private Cloud Base cluster uses HTTPS.
 - e) Select the Register KNOX endpoint (Optional) option.
 - f) KNOX IP Address - Enter the IP address of the Knox host for the CDP Private Cloud Base cluster.
 - g) KNOX Port - Enter the port for the Knox service.
 - h) Click CONNECT.

The Management Console acquires the configuration details from Cluster Connectivity Manager (CCM) service.

5. On the Classic Clusters page, click Files in the Step 2 pane.
6. In the Setup Connectivity Client dialog box, click `ccm-autossh-client`, `ssh_tunnel_setup_files` or `cluster_connectivity_setup_files`, and `ssh_tunnel_setup_files_knox` or `cluster_connectivity_setup_files_knox` links to download the RPM for AutoSSH from the specified location and the connectivity-installation scripts from CDP and Knox to your machine.
7. In the Setup Connectivity Client dialog box, click Close.
8. When the Cloudera Manager and Knox service are on the same host, perform the following steps to create two reverseSSH tunnels to the Cloudera Manager and Knox service:
 - a) In the command line interface, copy the RPM and ZIP files to the Cloudera Manager host.
 - b) To establish the first SSH tunnel between the Cloudera Manager and Management Console, SSH to the Cloudera Manager host.
 - c) Run the following command to install AutoSSH RPM using the `ccm-autossh-client` rpm file:

```
yum --nogpgcheck localinstall [***downloaded ccm_autossh-client-rpm***]
```

If the installation of CCM `autossh-client` rpm fails with “No package `autossh` available”, refer to [Troubleshooting: Install `autossh` package](#).


- d) Extract the `ssh_tunnel_setup_files.zip` or `cluster_connectivity_setup_files.zip` file to a temporary directory.
- e) Run the install script in the directory using the `./install.sh` command.
- f) Optionally, you can check the Cluster Connectivity Manager (CCM) service status using the `systemctl status ccm-tunnel@CM.service` command to verify whether the SSH tunnel is established.
- g) To establish the second SSH tunnel between the Knox service and Management Console, extract the `ssh_tunnel_setup_files_knox.zip` or `cluster_connectivity_setup_files_knox.zip` file to a temporary directory.
- h) Run the install script in the temporary directory using the `./install.sh` command.
- i) Enter the Knox topology details that you created earlier when the Enter Knox topology name that contains LDAP setup: prompt appears.
- j) Optionally, you can check the Cluster Connectivity Manager (CCM) service status using the `systemctl status ccm-tunnel@KNOX.service` command to verify whether the SSH tunnel is established.



Note: If you regenerate the script files, you cannot use the previously downloaded `ssh_tunnel_setup_files.zip` or `cluster_connectivity_setup_files.zip` file because the file is no longer valid.

9. When the Cloudera Manager and Knox service are on the different hosts, perform the following steps to create a SSH tunnel to Cloudera Manager and another SSH tunnel to the Knox service:
 - a) In the command line interface, copy the `ccm-autossh-client`, and `ssh_tunnel_setup_files` or `cluster_connectivity_setup_files` to the Cloudera Manager host.
 - b) To establish the first SSH tunnel between the Cloudera Manager host and and Management Console, SSH to the Cloudera Manager host.
 - c) Run the following command to install AutoSSH RPM using the `ccm-autossh-client rpm` file:


```
yum --nogpgcheck localinstall [***downloaded ccm_autossh-client-rpm***]
```
 - d) Extract the `ssh_tunnel_setup_files.zip` or `cluster_connectivity_setup_files.zip` file to a temporary directory.
 - e) Run the install script in the directory using the `./install.sh` command.
 - f) Optionally, you can check the Cluster Connectivity Manager (CCM) service status using the `systemctl status ccm-tunnel@CM.service` command to verify whether the SSH tunnel is established.
 - g) To establish the second SSH tunnel between the Knox service and Management Console, copy the `ccm-autossh-client` and `ssh_tunnel_setup_files_knox` or `cluster_connectivity_setup_files_knox` files to the Knox host.
 - h) SSH to the Knox host.
 - i) Run the following command to install AutoSSH RPM using the `ccm-autossh-client rpm` file:


```
yum --nogpgcheck localinstall [***downloaded ccm_autossh-client-rpm file***]
```
 - j) Extract the `ssh_tunnel_setup_files_knox.zip` or `cluster_connectivity_setup_files_knox.zip` file to a temporary directory.
 - k) Run the install script in the temporary directory using the `./install.sh` command.
 - l) Enter the Knox topology details that you created earlier when the Enter Knox topology name that contains LDAP setup: prompt appears.
 - m) Optionally, you can check the Cluster Connectivity Manager (CCM) service status using the `systemctl status ccm-tunnel@KNOX.service` command to verify whether the SSH tunnel is established.
-  **Note:** If you regenerate the script files, you cannot use the previously downloaded `ssh_tunnel_setup_files.zip` or `cluster_connectivity_setup_files.zip` file because the file is no longer valid.
10. On the Classic Clusters page, click Test Connection in the Step 2 pane to verify whether the tunnel connection is successful.
11. On the Classic Clusters page, click Register in the Step 3 pane.
12. In the Cluster Details dialog box, enter the Cloudera Manager credentials that has Admin access to Cloudera Manager and the cluster services.
13. Click CONNECT.
14. To complete the registration, enter the following details on the Classic Clusters page:
 - a) Cluster Location - Enter the geographical location of the Data Lake.
 - b) Data Center- Ensure that the data center name is the name that you provided for CDP Private Cloud Base cluster during registration.
 - c) Tags - Optionally, enter the tags for the cluster, if any.
 - d) Description- Optionally, enter a description.
15. Click Add.

Results

You can use the registered classic cluster in Replication Manager and Data Catalog.

Using classic clusters with a non-transparent proxy (CCMv1)

Navigation title: Non-transparent proxy

If your organization has a non-transparent proxy on the CM node, the following steps must be performed prior to classic cluster registration.



Note:

These steps only apply if you have a non-transparent proxy. You do not need to perform them if you have a transparent proxy.

When you register a cluster in CDP as a classic cluster, CDP installs CCM on the CM/Ambari node of CDH and HDP clusters to create a reverse tunnel, allowing communication with the CDP Control Plane to kick off replication jobs on schedule. To do this, CCM must be able to connect to the outside of the Data Center.

Steps

1. Download and install corkscrew on the CM node.

Corkscrew is an open source project with GPL licensing. You can download corkscrew from http://ftp.altlinux.org/pub/distributions/ALTLinux/Sisyphus/x86_64/RPMS.classic/corkscrew-2.0-alt1.qa1.x86_64.rpm.

For more information about corkscrew, see <https://github.com/bryanpkc/corkscrew/>.

2. CCM runs as root, so for the root user, create the .ssh directory and the following config files in the .ssh directory.

Note the following:

- Proxy server hostname does not need to contain http/https. Just use the hostname/IP address.
 - Proxy server port should be 443/80/8080 or similar.
 - Typical proxy_auth file would have “manishm:mypassword”.
- a. On CM node, create a /root/.ssh/config file with the following content, replacing the <PROXY-SERVER-HOSTNAME-IP-ADDRESS> and <PROXY-SERVER-PORT> with actual values:


```
ProxyCommand /usr/bin/corkscrew <PROXY-SERVER-HOSTNAME-IP-ADDRESS> <PROXY-SERVER-PORT> %h %p /root/.ssh/proxy_auth
```
 - b. On CM node, create a /root/.ssh/proxy_auth file with the following content, replacing the <PROXY-USERNAME> and <PROXY-PASSWORD> with actual values:


```
<PROXY-USERNAME>:<PROXY-PASSWORD>
```
 - c. Set ownership of the proxy_auth file as follows:


```
chmod 600 /root/.ssh/proxy_auth
```
 - d. If the proxy server does not need user id and password, then set just this in the /root/.ssh/config file as follows, replacing the <PROXY-SERVER-HOSTNAME-IP-ADDRESS> and <PROXY-SERVER-PORT> with actual values:


```
ProxyCommand /usr/bin/corkscrew <PROXY-SERVER-HOSTNAME-IP-ADDRESS> <PROXY-SERVER-PORT>
```

Once you’ve performed these steps, you can proceed to registering your cluster in CDP.

Troubleshooting classic cluster registration errors (CCMv1)

Navigation title: Troubleshooting classic clusters

While trying to resume the registration process, you can identify the problem behind the error and fix it.

Issues during registration in CDP

Error or issue details	Resolution
Alert: Registration is pending for a cluster with the same details.	<p>A cluster with the same IP Address and Datacenter cannot be registered again.</p> <p>Check if you have registered this cluster already. To check this, navigate to the Classic Clusters page and search for your cluster in the list.</p> <p>Check if the IP address is correct.</p> <p>Provide a different DataCenter name.</p>

Cluster side issues

Error or issue details	Resolution
When installing the AutoSSH rpm on a cluster node, fetching the ccm-autossh-client rpm package failed. Multiple mirrors were tried but a lot of them resulted in 404/ timed out errors.	Install autossh independently and then try installing the script.
Connection refused even though the systemctl status ccm-tunnel@CM.service shows that the autossh client is running.	Make sure you copied the right ssh setup files for the cluster or check if the port number CCM_TUNNEL_SERVICE_PORT in cm_reverse_tunnel.conf or cluster_connectivity.conf is your Cloudera Manager's port number.
When running install script on cluster node, the ssh tunnel could not be established (as indicated in the logs generated using journalctl -xe).	Check if outbound connection to CDP or HDP control plane (host/port = CCM_HOST/CCM_SSH_PORT) is allowed. Make sure to check the firewall rules.
Test connection failure	<p>Try the following:</p> <ul style="list-style-type: none"> Check if the reverse SSH tunnel is running on the Knox node: <pre>systemctl status ccm-tunnel@KNOX.service</pre> If the tunnel is not running, start the tunnel by running the install script. If the tunnel is active, check if there are any Ranger policies set up to deny access to the cluster. If such policies exist on the cluster, modify or set up policies to allow access to the cluster cdp_default topology. If the tunnel is active, check if the port number entered during registration is correct. If the port number is incorrect, delete the registration attempt from the UI, remove all the setup-related files from the cluster node and re-register the cluster with the correct information. If the port number is correct, check if outbound connection to CDP control plane (host/port = CCM_HOST/CCM_SSH_PORT) is allowed on the directory which has the ssh setup files. <ul style="list-style-type: none"> cat reverse_tunnel.conf or cat cluster_connectivity.conf and collect the following properties - CCM_HOST, CCM_SSH_PORT, CCM_TUNNEL_INITIATOR_ID Check if the NLB is reachable from the node: <ul style="list-style-type: none"> On the node, execute the command nslookup <CCM_HOST> to check if the NLB is reachable from the node. If this fails, then the traffic to the cloudera network is blocked on the customer's VPC. The customer should check the outbound rules on their VPC to make sure that the traffic to the cloudera network is allowed.

CCM issues

Classic Clusters registration uses CCM, enabled by the installation of the CCM client and secrets on the CM node.

Error or issue details	Resolution
Unable to get endpoint from CCM for key <key>	<p>The Network Load Balancers (NLBs) might be unreachable.</p> <ol style="list-style-type: none"> 1. To verify, find the autossh process on the FreeIPA master host by issuing the <code>ps aux grep autossh</code> command. 2. Modify the command by adding the process name: <pre>ssh -o ConnectTimeout=30 -o ServerAliveInterval=30 -o ServerAliveCountMax=3 -o UserKnownHostsFile=/etc/ccm/ccm.pub ... -vvv</pre> 3. Outgoing traffic should be allowed in the port range 6000-6049 on the legacy CDH/HDP cluster. 4. If there is a proxy in the network it is also worth checking the ssh proxy configurations under: <ul style="list-style-type: none"> • /root/.ssh/config • /root/.ssh/proxy_auth

Other issues

Error or issue details	Resolution
If Cluster name and Display name are different, few details are missing from the cluster detail page.	Cluster name and Display name must be the same.

Troubleshooting: Install autossh package

During step 2 of the classic cluster registration process, you install an rpm to establish connectivity between the cluster and CDP Control Plane. If the installation fails, you can install the autossh using the following steps.

During step 2 of the classic cluster registration process, you install an rpm to establish connectivity between the cluster and CDP Control Plane. The CCMv1 autossh client rpm has dependency on the autossh package which is used for creating the reverse SSH tunnel from the cluster node (Knox/CM) to CDP Control Plane. The instructions on the screen instruct you to run the following command:

```
sudo yum --nogpgcheck localinstall ccm-autossh-client-0.2-20201211113132gita7b10d3.x86_64.rpm
```

If the installation of CCM autossh-client rpm fails with “No package autossh available”, you must download and install the autossh package on the node where you are setting up connectivity.

Steps

1. Run these commands to install the autossh package:

```
yum install wget
wget https://download-ib01.fedoraproject.org/pub/epel/7/x86_64/Packages/a/autossh-1.4g-1.el7.x86_64.rpm
sudo yum --nogpgcheck localinstall autossh-1.4g-1.el7.x86_64.rpm
```

2. Next, run the following command to install autossh-client:

```
sudo yum --nogpgcheck localinstall ccm-autossh-client-0.2-20201211113132gita7b10d3.x86_64.rpm
```

Adding classic clusters with CCMv2

If your CDP tenant has been granted the CDP_CCM_V2 entitlement to use CCMv2, refer to the following documentation for instructions on how to register your existing on-premise cluster in CDP Public Cloud.

Adding an HDP cluster (CCMv2)

Navigation title: Adding an HDP cluster

You must register HDP clusters with CDP before you can use them with CDP Public Cloud services and components.

About this task

To ensure optimum security, clusters within the customer environment are not accessible for communication: They have private IP addresses and cannot be accessed outside the firewall. To add your cluster to CDP, a communication line needs to be established.

The CCM inverted proxy solves the problem by establishing a connection from the on-premise cluster to CDP. You must download and install the jumpgate agent and the connectivity install scripts to establish a secure two-way communication channel. The jumpgate agent ensures that the connectivity is stable. The connectivity scripts and their installation ensure safe connectivity and communication.

Before you begin

- HDP clusters must be managed by Ambari. HDP clusters that are not managed by Ambari cannot be registered to CDP.
- HDP clusters must include Knox.
- HDP clusters must include Ranger policy settings
- LDAP/AD must be set up and synced in Ambari.

LDAP settings are automatically detected from the default topology setup in Knox. If the default topology does not have the LDAP setup, you will be asked to provide another topology name where you have configured the LDAP. If that topology has LDAP, the setup continues. If the LDAP is not configured, you will receive an error message.

- Kerberos must be enabled on the HDP cluster and the LDAP/AD must be set up in the Kerberos authentication so that the same set of LDAP/AD credentials can be used to access Ambari APIs as well as Beacon APIs.
- All clusters must meet the requirements identified in [Prerequisites for adding classic clusters](#).



Caution: After you register an HDP cluster in CDP, do not change the cluster name in Ambari. A cluster name change in Ambari does not currently propagate to CDP, which can result in issues when using the HDP cluster with CDP clusters and components.



Caution: The jumpghost and connectivity-install scripts files must be stored in a secure environment

Steps

The process to register an HDP cluster is as follows:

1. Log in to CDP and navigate to the Management Console.
2. Click Classic Clusters in the left navigation panel.
3. Click Add Cluster.

CDP displays the Cluster Details dialog box.

4. If you are a first time user, under Step 1 in the Register Classic Cluster wizard, click GET STARTED. If you are not a first time user, click the ADD CLUSTER button on the right side of the listing page.

CDP displays the Cluster Details dialog box.

5. Select HDP.
6. Provide the following connectivity information for your new cluster:
 - a. Knox IP Address
 - b. Knox port
 - c. Data center
 - d. Click CONNECT

Step 1 might take up to 5 minutes if you are adding the cluster for the first time. After Classic Cluster establishes connection, CDP will highlight Step 2 in the Register Classic Cluster wizard.

7. Start the download and installation process for the connectivity files by clicking the Files button in Step 2 of the wizard.
8. Follow the instructions in the Setup Connectivity Client dialog box. You need to download the jumpgate-agent rpm file and the cluster_connectivity_setup_files zip files and copy them to your Knox node. Or if your Knox is running in HA mode, you need to copy the files to the Knox proxy host in the cluster:
 - a. Download the jumpgate-agent RPM and cluster_connectivity_setup_files.zip.
 - b. In the command line interface, copy the two files to the Knox or Knox proxy host.
 - c. SSH to the host.
 - d. Install the jumpgate-agent rpm using:

```
yum --nogpgcheck localinstall < downloaded-jumpgate-agent-rpm >
```

- a. Unzip the cluster_connectivity_setup_files file. Inside this zip file there is a script install.sh.
- b. Run install.sh by using ./install.sh command.
- c. Check service status to see if the agent has been connected:

```
systemctl status jumpgate-agent.service
```

9. Enter the following information as the install script prompts for it:
 - a. Enter Ambari URL (http(s)://host:[port]):
 - b. Enter Ambari Username:
 - c. Enter Ambari Password:
10. If Knox is not installed on a proxy server, proceed to Step 12. Classic Cluster sets up the topology for the Knox server and establishes the connection.
11. If Knox is installed on a proxy server, Classic Cluster displays the following message: We discovered that your Knox is installed in HA mode. Please confirm if this node is your proxy node (yes/no): Enter yes. Classic Cluster generates XML content that you will need to add to your Knox hosts. Classic Cluster also displays three steps you must perform on all of your Knox hosts:
 - a. Copy the generated XML to /usr/hdp/current/knox-server/conf/topologies/cdp_default.xml
 - b. Run `chown knox:knox cdp_default.xml`
 - c. Check the Knox logs/deployment directory to verify that the cdp_default topology is deployed.
 - d. After you have completed steps a through c on all of your Knox hosts, type Enter to continue. This sets up the topology for the Knox server or Knox proxy host and establishes the connection.
12. On the Classic Clusters page, click Test Connection in the Step 2 pane to verify whether the connection is successful.

13. CDP starts checking the connectivity with the HDP cluster. When the connectivity is successful, proceed to Step 3 in the wizard.

If the connection attempts fail or if there is an error in the connectivity, CDP displays troubleshooting information in Step 2 of the registration wizard. Follow the troubleshooting information to fix the connectivity error, then click Test connection.



Note: After you download the files, the cluster_connectivity_setup files download is disabled. At this point, you can regenerate the cluster connectivity setup files using the regenerate files option. This option comes in handy if you lose the files previously downloaded before you can set the cluster connectivity in the cluster.

14. Click Register in Step 3 of the registration wizard.

15. In the Cluster Details dialog box, provide the username and password to access the cluster, then click CONNECT.

The user should have admin access to the cluster services.

16. Finish registering the cluster by providing the following information.

- a. Cluster Location
- b. Data Center
- c. Tags (optional)
- d. Description (optional)

If LDAP is not set up on the default topology, the system will ask for the following additional information: Enter knox topology name that contains LDAP setup.

17. Click ADD.

Result

Once the cluster has been registered, you can use it with CDP.

Adding a CDH cluster (CCMv2)

Navigation title: Adding a CDH cluster

You must register CDH clusters in CDP before you can use them with CDP Public Cloud services and components.

About this task

To ensure optimal security, clusters within the customer environment are not accessible for communication; They have private IP addresses and cannot be accessed outside the firewall. To add your cluster to the CDP, a communication line needs to be established.

The CCM inverted proxy solves the problem by establishing a connection from the on-premise cluster to CDP. You must download and install the jumpgate agent and the connectivity install scripts to establish a secure two-way communication channel. The jumpgate agent ensures that the connectivity is stable. The connectivity scripts and their installation ensure safe connectivity and communication.

The high-level steps to register a CDH cluster using CCM are as follows:

1. In the CDP Management Console, you enter the private IP address of your cluster and provide the cluster details.
2. You download the jumpgate agent rpm from the specified location and the connectivity installation scripts from CDP on to the cluster.
3. You install jumpgate agent on the cluster.
4. You register the cluster for performing further operations.

Detailed steps are provided below.

Before you begin

- CDH clusters must have been created using Cloudera Manager. Clusters that are not managed by Cloudera Manager cannot be registered to CDP.
- All clusters must meet the requirements identified in [Prerequisites for adding classic clusters](#).



Caution: After you register a CDH cluster in CDP, do not change the cluster name in Cloudera Manager. A cluster name change in Cloudera Manager does not currently propagate to CDP, which can result in issues when using the CDH cluster with CDP clusters and components.



Caution: The jumpgate agent and connectivity-install scripts files must be stored in a secure location.

Steps

Perform the following steps to add a CDH cluster:

1. Log in to CDP and navigate to the Management Console.
2. Click Classic Clusters in the left navigation panel.
3. Click Add Cluster.

CDP displays the Cluster Details dialog box.

4. Click CDH.
5. Provide the connectivity information for your new cluster:



Note: Make sure that the Data Center name is different from the Data Center names that have already been registered. If the Data Center name exists, make sure that the combination of the Data Center name and the cluster name is unique. Else, you may get an error when you try to add a cluster with an existing Data Center-cluster name combination.

- a. Cloudera Manger IP address
- b. Cloudera Manager Port
- c. Data center
- d. Select the My cluster runs on HTTPS option if the CDH cluster uses HTTPS.

Step 1 might take up to five minutes if you are adding the cluster for the first time. After Step 1 is complete, CDP will highlight Step 2.

6. Start the download and installation process for the connectivity files by clicking the Files button in the Step 2 pane.
7. Follow the instructions in the Setup Connectivity Client dialog box. You need to download the jumpgate-agent rpm file and the cluster_connectivity_setup_files zip file onto Cloudera Manager host in your new cluster:
 - a. Download the jumpgate-agent RPM and cluster_connectivity_setup_files.
 - b. In the command line interface, copy the two files to the Cloudera Manager host.
 - c. SSH to the Cloudera Manager host.
 - d. Install the jumpgate-agent rpm using:

```
yum --nogpgcheck localinstall <
```

```
downloaded-jumpgate-agent-rpm >
```

- e. Unzip the cluster_connectivity_setup_files file. Inside this zip file there is a script install.sh.
- f. Run install.sh by using ./install.sh command.
- g. Check service status to see if the agent has been connected:

```
systemctl status jumpgate-agent.service
```

On the Classic Clusters page, click Test Connection in the Step 2 pane to verify whether the connection is successful.

After CDP successfully connects to your new cluster, it will highlight Step 2. the connectivity with the cluster. When the connectivity is successful, proceed to Step 3 in the UI.

If the connection attempts fail or if there is an error in the connectivity, CDP displays troubleshooting information in the Step 2 pane. Follow the troubleshooting information to fix the connectivity error, then click Test connection.



Note: After you download the files, the cluster_connectivity_setup file download is disabled. At this point, you can regenerate the cluster connectivity setup files using the regenerate files option. This option comes in handy if you lose the files previously downloaded before you can set the cluster connectivity in the cluster.

8. Click Register in the Step 3 pane.
9. Provide the username and password of the Cloudera Manager user to access the cluster.
10. Finishing registering the cluster by providing the following information:
 - a. Cluster Location
 - b. Data Center
 - c. Tags (optional)
 - d. Description (optional)
11. Click Submit.

Result

Once the cluster has been registered, you can use it with CDP.

Adding a CDP Private Cloud Base cluster (CCMv2)

Navigation title: Adding a CDP PvC Base cluster

You must register CDP Private Cloud Base clusters with CDP before you can use them with CDP Public Cloud services and components.

To ensure optimum security, clusters within the customer environment are not accessible for communication: They have private IP addresses and cannot be accessed outside the firewall. To add your cluster to the CDP, a communication line needs to be established.

The CCM inverted proxy solves the problem by establishing a connection from the on-premise cluster to CDP. You must download and install the jumpgate agent and the connectivity install scripts to establish a secure two-way communication channel. The jumpgate agent ensures that the connectivity is stable. The connectivity scripts and their installation ensure safe connectivity and communication.

Note: The jumphost and connectivity-install scripts files must be stored in a secure environment.

To register the CDP Private Cloud Base cluster as a classic cluster, you enter the CDP Private Cloud Base cluster details. The Management Console acquires the configuration details from Cluster Connectivity Manager (CCM) service and saves it as ZIP files. You download the ZIP files, install the acquired configurations, and then register the CDP Private Cloud Base cluster as a classic cluster.

Note: After you register a CDP Private Cloud Base cluster in CDP, do not change the cluster name in Cloudera Manager. A cluster name change in Cloudera Manager does not currently propagate to CDP, which can result in issues when using clusters with CDP clusters and components.

All the clusters must meet the requirements identified in [Prerequisites for adding classic clusters](#).

You have two options for registering your CDP Private Cloud Base cluster:

- If you would like to use the CDP Private Cloud Base cluster with Replication Manager, register the cluster using Cloudera Manager.
- If you would like to use the CDP Private Cloud Base cluster with Replication Manager and Data Catalog, register the cluster using Cloudera Manager and Knox.



Important: CDP Private Cloud Base clusters can be used in Data Catalog by registering them using Cloudera Manager and Knox endpoints. Note that this is a technical preview feature and is under development. Do not use this in your production environment. If you have feedback, contact Support by logging a case on the Cloudera Support Portal at <https://my.cloudera.com/support.html>. Technical preview features are not guaranteed troubleshooting and fixes

For CDP Private Cloud Base cluster registration steps in CDP, see the following documentation:

Adding CDP Private Cloud Base cluster for use in Replication Manager (CCMv2)

Navigation title: Replication Manager use case

Register a CDP Private Cloud Base cluster as a classic cluster using Cloudera Manager so that you can use this cluster as a source cluster in Replication Manager.

Before you begin

All the clusters must meet the requirements identified in [Prerequisites for adding classic clusters](#).

Steps

1. Log in to CDP Management Console.
2. Click Classic Clusters.
3. On the Classic Clusters page, click ADD CLUSTER.
4. In the Add Cluster dialog box, navigate to the CDP Private Cloud Base tab and enter the following details:
 - a. If your cluster is not reachable by a public network, click “My cluster is accessible only in my private network”.
 - b. Cloudera Manager IP address - Enter the IP address of the Cloudera Manager of the CDP Private Cloud Base cluster. The Management Console uses this IP address to identify the cluster for registration purposes.
 - c. Cloudera Manager Port - Enter the port of the Cloudera Manager of the CDP Private Cloud Base cluster.
 - d. Data center - Enter a unique data center name for the CDP Private Cloud Base cluster.
 - e. Select the My cluster runs on HTTPS option if the CDP Private Cloud Base cluster uses HTTPS.
 - f. Clear the Register KNOX endpoint (Optional) option, if selected.
 - g. Click CONNECT.

The Management Console acquires the configuration details from Cluster Connectivity Manager (CCM) service. After CDP successfully connects to your new cluster (which should take no more than 5 minutes), it will highlight Step 2.

5. On the Classic Clusters page, click Files in the Step 2 pane.
6. Follow the instructions in the Setup Connectivity Client dialog box. You need to download the jumpgate-agent rpm file and the cluster_connectivity_setup_files zip file onto Cloudera Manager host in your new cluster:
 - a. In the command line interface, copy the jumpgate-agent RPM and cluster_connectivity_setup_files.zip to the Cloudera Manager host.
 - b. SSH to the Cloudera Manager host.
 - c. Install the jumpgate-agent rpm using:

```
yum --nogpgcheck localinstall <
```

```
downloaded-jumpgate-agent-rpm >
```

- a. Unzip the cluster_connectivity_setup_files file. Inside this zip file there is a script install.sh.
- b. Run install.sh by using ./install.sh command.
- c. Check service status to see if the agent has been connected:

```
systemctl status jumpgate-agent.service
```



Note: If you regenerate the script files, you cannot use the previously downloaded cluster_connectivity_setup_files.zip file because the file is no longer valid.

7. On the Classic Clusters page, click Test Connection in the Step 2 pane to verify whether the connection is successful.
8. Click Register in the Step 3 pane.
9. In the Cluster Details dialog box, enter the Cloudera Manager credentials that have Admin access to Cloudera Manager and the cluster services.
10. Click CONNECT.
11. To complete the registration, enter the following details on the Classic Clusters page:
 - a. Cluster Location - Enter the geographical location of the Data Lake.
 - b. Data Center - Ensure that the data center name is the name that you provided for CDP Private Cloud Base cluster during registration.
 - c. Tags - Optionally, enter the tags for the cluster.
 - d. Description - Optionally, enter a description.
12. Click Add.

Result

You can use the registered classic cluster in the Replication Manager.

Adding CDP Private Cloud Base cluster for use in Replication Manager and Data Catalog (CCMv2)

Navigation title: Replication Manager and Data Catalog use case

Register a CDP Private Cloud Base cluster as a classic cluster using Cloudera Manager and Knox endpoints so that you can use this cluster in Replication Manager and Data Catalog.

Before you begin

All the clusters must meet the requirements identified in [Prerequisites for adding classic clusters](#).

Additionally, ensure that the following components and roles are available:

- The CDP Private Cloud Base cluster has an active Knox service.
- You can proxy to Cloudera Manager through Knox for communication purposes.

Disposition: / Status:

Removed link to Knox UI doc that was removed.

- LDAP is configured in the Cloudera Manager of CDP Private Cloud Base cluster. For more information, see [Configure authentication using an LDAP-compliant identity service](#).
- A minimum of one LDAP user with the Full Administrator role.
- An LDAP-based topology cdp_default.xml with CM-API, CM-UI, ATLAS, ATLAS-API, RANGERUI, and RANGER services exists. The topology name is used during the classic cluster registration process.

Disposition: / Status:

Removed link to Knox UI doc that was removed.



Note: If there are policies that restrict access through Knox, then add the topology name to the cdp_default Ranger policy so that the Ranger policies can communicate through Knox.



Important: CDP Private Cloud Base clusters can be used in Data Catalog by registering them using Cloudera Manager and Knox endpoints. Note that this is a technical preview feature and is under development. Do not use this in your production environment. If you have feedback, contact Support by logging a case on the Cloudera Support Portal at <https://my.cloudera.com/support.html>. Technical preview features are not guaranteed troubleshooting and fixes.

Steps

1. Log in to CDP Management Console.
2. Click Classic Clusters.
3. On the Classic Clusters page, click ADD CLUSTER.
4. In the Add Cluster dialog box, navigate to the CDP Private Cloud Base tab and enter the following details:
 - a. If your cluster is not reachable by a public network, click “My cluster is accessible only in my private network”.
 - b. Cloudera Manager IP address - Enter the IP address of the Cloudera Manager of the CDP Private Cloud Base cluster. The Management Console uses this IP address to identify the cluster for registration purposes.
 - c. Cloudera Manager Port - Enter the port of the Cloudera Manager of the CDP Private Cloud Base cluster.
 - d. Data center - Enter a unique datacenter name for the CDP Private Cloud Base cluster.
 - e. Select the My cluster runs on HTTPS option if the CDP Private Cloud Base cluster uses HTTPS.
 - f. Select the Register KNOX endpoint (Optional) option.
 - g. KNOX IP Address - Enter the IP address of the Knox host for the CDP Private Cloud Base cluster.
 - h. KNOX Port - Enter the port for the Knox service.
 - i. Click CONNECT.

The Management Console acquires the configuration details from Cluster Connectivity Manager (CCM) service. After CDP successfully connects to your new cluster (which should take no more than 5 minutes), it will highlight Step 2.

5. On the Classic Clusters page, click Files in the Step 2 pane.
6. Follow the instructions in the Setup Connectivity Client dialog box. You need to download the jumpgate-agent rpm file and the cluster_connectivity_setup_files zip file onto Cloudera Manager host in your new cluster:
 - a. In the command line interface, copy the RPM and ZIP files to the Cloudera Manager host.
 - b. SSH to the Cloudera Manager host.
 - c. Install the jumpgate-agent rpm using `yum --nogpgcheck localinstall < downloaded-jumpgate-agent-rpm >`
 - d. Unzip the cluster_connectivity_setup_files file. Inside this zip file there is a script install.sh.
 - e. Run install.sh by using `./install.sh` command.
 - f. Check service status to see if the agent has been connected: `systemctl status jumpgate-agent.service`



Note: If you regenerate the script files, you cannot use the previously downloaded cluster_connectivity_setup_files.zip file because the file is no longer valid.

7. On the Classic Clusters page, click Test Connection in the Step 2 pane to verify whether the connection is successful.
8. On the Classic Clusters page, click Register in the Step 3 pane.
9. In the Cluster Details dialog box, enter the Cloudera Manager credentials that have Admin access to Cloudera Manager and the cluster services.
10. Click CONNECT.
11. To complete the registration, enter the following details on the Classic Clusters page:
 - a. Cluster Location - Enter the geographical location of the Data Lake.
 - b. Data Center - Ensure that the data center name is the name that you provided for CDP Private Cloud Base cluster during registration.
 - c. Tags - Optionally, enter the tags for the cluster, if any.
 - d. Description - Optionally, enter a description.
12. Click Add.

Result

You can use the registered classic cluster in Replication Manager and Data Catalog.

Using classic clusters with a non-transparent proxy (CCMv2)

Navigation title: Non-transparent proxy

If your organization has a non-transparent proxy on the CM/Knox node, the following steps must be performed prior to classic cluster registration.



Note: These steps only apply if you have a non-transparent proxy. You do not need to perform them if you have a transparent proxy.



Note: An https proxy is supported only if the certificate is added to the system trust store.

When you register a cluster in CDP as a classic cluster, CDP installs CCM on the CM/Knox node of CDH and HDP clusters to establish connection between the on-premise cluster and CDP, allowing communication with the CDP Control Plane to kick off replication jobs on schedule. To do this, CCM must be able to connect to the outside of the Data Center.

Steps

Create the following file on the Cloudera Manager node (in case of a CDH or CDP Private Cloud Base cluster) or on the Knox node (in case of an HDP cluster):

```
/etc/cdp/proxy.env
```

The file should include a proxy link. The format of the file should be:

```
https_proxy=http://<username>:<password>@<proxy.com>
```

- The <username> and <password> should be replaced with an actual username that allows access to the proxy.
- The <proxy.com> should be replaced with the URL of the proxy server.

Once you've performed these steps, you can proceed to registering your cluster in CDP.

Troubleshooting classic cluster registration errors (CCMv2)

Navigation title: Troubleshooting classic clusters

While trying to resume the registration process, you can identify the problem behind the error and fix it.

Issues during registration in CDP

Error or issue details	Resolution
Alert: Registration is pending for a cluster with the same details.	<ul style="list-style-type: none"> • A cluster with the same IP Address and Data Center cannot be registered again. • Check if you have registered this cluster already. To check this, navigate to the Classic Clusters page and search for your cluster. • Check if the IP address is correct. • Provide a different Data Center name.

Cluster side issues

Error or issue details	Resolution
Connection refused even though the <code>systemctl status jumpgate-agent.service</code> shows that the jumpgate agent is running.	Make sure that you copied the right setup files for the cluster or check if the port number <code>CCM_TUNNEL_SERVICE_PORT</code> in <code>cluster_connectivity.conf</code> is your Cloudera Manager's port number in case of CDH/CDP Private Cloud Base cluster and Knox port number in case of HDP cluster.
Test connection failure	<p>Try the following:</p> <ul style="list-style-type: none"> Check if the jumpgate agent is running: <code>systemctl status jumpgate-agent.service</code> If the jumpgate agent is not running, start the agent by running the install script. If the agent is active, check if the port number entered during registration is correct. If the port number is incorrect, delete the registration attempt from the UI, remove all the setup-related files from the cluster node and re-register the cluster with the correct information. If the port number is correct, check if outbound connection to the CDP Control Plane is allowed. <ul style="list-style-type: none"> <code>cat cluster_connectivity.conf</code> and collect the <code>RELAY_SERVER</code> Check if the <code>RELAY_SERVER</code> is reachable from the node: <ul style="list-style-type: none"> On the node, execute the command <code>nslookup <RELAY_SERVER></code> to check if the <code>RELAY_SERVER</code> is reachable from the node. Or execute curl command as mentioned here for connectivity check: <code>curl https://<RELAY_SERVER>:443</code> Or execute telnet command as mentioned here: <code>telnet <RELAY_SERVER> 443</code> If this fails, then the traffic to the cloudera network is blocked on the customer's VPC. You should check the outbound rules on your VPC to make sure that the traffic to the Cloudera network is allowed. If the agent is active, check if there are any Ranger policies set up to deny access to the cluster. If such policies exist on the cluster, modify or set up policies to allow access to the cluster <code>cdp_default topology</code> If the agent is active, check if there are any Ranger policies set up to deny access to the cluster. If such policies exist on the cluster, modify or set up policies to allow access to the cluster <code>cdp_default topology</code>

CCM issues

Classic Clusters registration uses CCM, enabled by the installation of the CCM client and secrets on the CM node.

Error or issue details	Resolution
<p>"Unable to identify a backendId for this request. This usually happens if either the backendId is invalid or the agent has not yet connected to the relay server"</p> <p>OR</p> <p>"No route found for backend with id <backend-id>. This usually happens if either the backendId is invalid or the agent has not yet connected to the relay server."</p>	<p>The CCMv2 Network Load Balancers (NLBs) might be unreachable. To verify, execute <code>cat cluster_connectivity.conf</code> and collect the <code>RELAY_SERVER</code> field's value and make sure that the <code>RELAY_SERVER</code> on port 443 is reachable from the legacy CDH/HDP master nodes.</p>

Other issues

Error or issue details	Resolution
If Cluster name and Display name are different, few details are missing from the cluster detail page.	Cluster name and Display name must be the same.

If registering your cluster for use with Data Catalog, also check the following:

- Check Proxy to Cloudera Manager through Apache Knox is enabled.
- Make sure that you created the `cdp_default` topology on the Knox host with required services and LDAP configuration; it's a prerequisite.

Disposition: / Status:

Removed link to Knox UI doc that was removed.

Resume the cluster registration

If the cluster registration does not succeed due to some factors, you can resume the registration, identify and fix any issues, and complete the registration process.

Procedure

1. Go to the Classic Clusters page.
2. From the list of unregistered clusters that appears, select your cluster and start the process from the point you left off.

Delete an unregistered classic cluster

If you started registering an on-prem cluster in CDP PC using CCMv1 and have not yet completed it, we recommend that you delete the unregistered cluster and re-register it using CCMv2.

Classic clusters that are already registered with CCMv1 continue to use CCMv1, but new cluster registrations should be using CCMv2. If you started registering an on-prem cluster in CDP PC using CCMv1 and have not yet completed it, we recommend that you delete the unregistered cluster and re-register it using CCMv2. We suggest this as it will save the future efforts required to upgrade from CCMv1 to CCMv2.

Steps

1. If you reached step 2 or 3 in classic cluster registration, you should first stop the CCM tunnel and then clean up the setup files that you must have downloaded for CCM v1. Run the following commands on the Cloudera Manager host (CDH or CDP PvC Base cluster) or on the KNOX host (HDP cluster).
 - a. Check the tunnel status:

```
systemctl status ccm-tunnel@CM.service
```

- b. Use below commands to stop tunnel and cleanup CCMv1 resources:

- CentOS7/RHEL7:

CDH or CDP PvC Base cluster

```
systemctl stop ccm-tunnel@CM.service
```

```
yum remove autossh
```

HDP cluster

```
systemctl stop ccm-tunnel@KNOX.service
yum remove autossh
```


- CentOS6/RHEL6:

CDH or CDP PvC Base cluster

```
service reverse-tunnel stop 'CM'
yum remove autossh
```

HDP cluster

```
service reverse-tunnel stop 'KNOX'
yum remove autossh
```

2. Delete the cluster from the classic cluster UI by clicking on the  context menu next to the cluster name and selecting Remove.
3. Enter the cluster name and click Remove.

Managing a classic cluster













You can manage a CDH, HDP, or CDP Private Cloud Base cluster using the Classic Cluster user interface.

About this task

You can perform the following management functions using the Classic Cluster user interface.

Procedure

1. You can delete a cluster if it has not been completely registered by clicking the delete icon next to the cluster entry in the list of clusters.

Type 	Status Registration In Progress	Name NA	IP Address 172.0.0.1	Data center test12345	Port 8989	Last Updated 21-Apr-2020 10:03:11	 
Type 	Status Registration In Progress	Name NA	IP Address 172.27.128.68	Data center test1170	Port 7180	Last Updated 16-Apr-2020 02:41:05	 
Type 	Status Registration In Progress	Name NA	IP Address 172.27.128.68	Data center test3625	Port 7180	Last Updated 14-Apr-2020 22:12:21	 
Type 	Status Registration In Progress	Name NA	IP Address 172.27.128.68	Data center test2485	Port 7180	Last Updated 14-Apr-2020 22:10:02	 

2. You can view cluster details from the UI by clicking the cluster name in the list of clusters.

Classic Clusters						
Type	Status	Name	IP Address	Data center	Port	Last updated
	Active	Cluster 1	172.27.82.4	\$77Multi2	7180	30-Apr-2020 02:25:21
	Active	gfi	172.27.126.193	hdp26	8080	30-Apr-2020 01:42:28
	Active	mycluster0	172.27.68.6	344HDPPrivateCCM	8443	29-Apr-2020 18:07:53

Classic Cluster displays the Details dialog box.

CLUSTER

Management Console

Dashboard

Environments

Data Lakes

User Management

Data Hub Clusters

Data Warehouses

ML Workspaces

Classic Clusters

Cost Management

Get Started

Help

Arun Sarin

Classic Clusters / Details

Cluster 1

Last updated 30 Apr 2020 22:25:21

LOCATION

DATA CENTER

CLUSTER VERSION

NODES

TAGS

REGISTERED AT

REGISTERED BY

Chennai, India

\$77Multi2

7.1.1

3/3

\$77

Wed Apr 29 2020

Arun Sarin

INFORMATION

Cluster Services

Description

577

Cluster Services

ZOOKEEPER

HDFS

HIVE

YARN

KAFKA

Connectivity

Reachable

Security Type

NA

No of Services

18

DataNodes

2 / 0

NodeManager Heap Size

131 MB / 4 GB

9.23%

HDFS Disk Space

4 GB / 484 GB

0.93%

NodeManagers

2 / 0

ResourceManager Heap Size

107 MB / 1027 MB

10.42%

ResourceManager Uptime

a month

Average Kafka Broker Topics In

0 / sec

Average Kafka Broker Topics Out

0 / sec

Average Kafka Broker Messages In

0

Kafka Active Controller Count

1

Kafka Brokers

3 / 0

Kafka Replica Max Lag

0

Kafka Partition Count

0

Kafka Replica Manager Leader Count

0

Kafka UnderReplicated Partitions Count

0

CLUSTER REGISTRATION DETAILS

Cluster Manager URL

http://172.27.82.4:7180

Registration Mode

Basic Auth

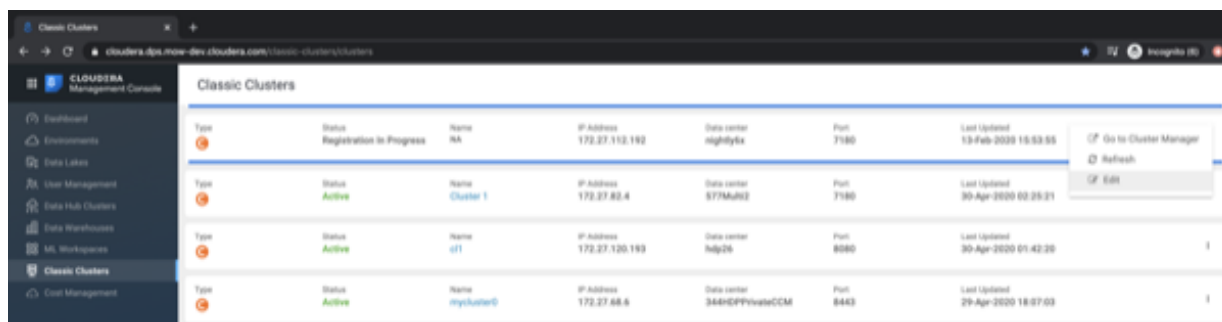
Knox URL

NA

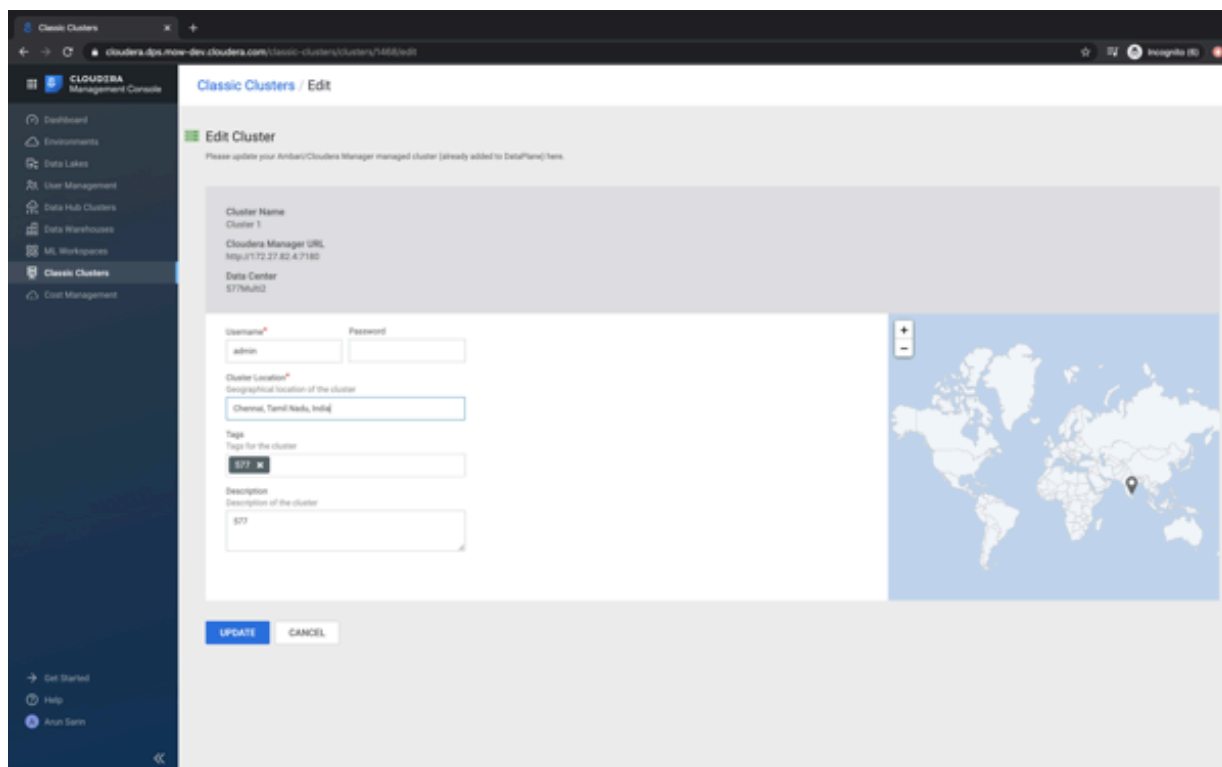
SSL Validation

DISABLED

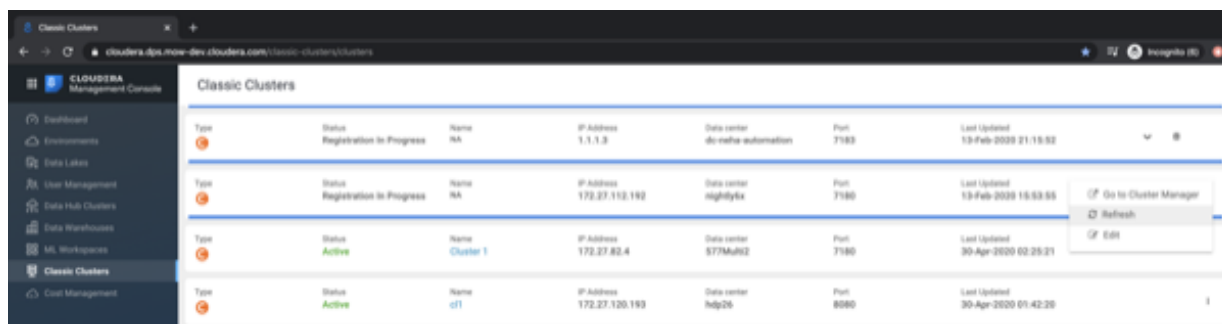
- You can edit a cluster from the UI by choosing the Edit option from the pull down menu next to the cluster entry in the list of clusters.



Classic Cluster displays the Edit dialog box.



- You can refresh cluster information for a registered cluster by choosing the Refresh option from the pull down menu next to the cluster entry in the list of clusters.



Upgrading a classic cluster from CCMv1 to CCMv2


Existing classic clusters (CDH, HDP, or CDP Private Cloud Base) registered in CDP Public Cloud's Management Console with Cluster Connectivity Manager v1 (CCMv1) should be upgraded to use CCMv2.


Before you begin

Prior to upgrading, open ports for CCMv2. If you would like to use Cluster Connectivity Manager v2 (CCMv2), ensure that outgoing traffic is allowed on port 443 on the legacy CDH, HDP, or CDP Private Cloud Base cluster.

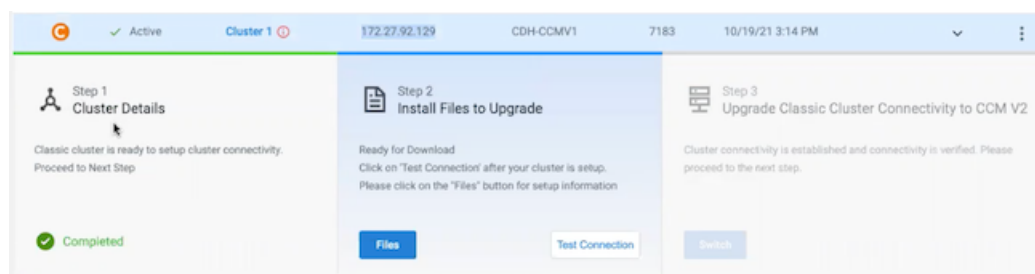
Procedure

1. Log in to the CDP web interface.
2. Navigate to the Management Console.
3. From the left navigation pane, select Classic clusters.
4. Find your previously registered classic cluster.
- 5.

If your cluster need to be upgraded, you will see the  icon next to its name. After clicking on this icon, you will see the Upgrade Now button, as shown in the following screenshot:



6. Click the Upgrade Now button to start the upgrade process.
7. Information about three upgrade steps is displayed:



8. CDP automatically starts step 1. It normally takes 2-4 minutes. Once step 1 is marked as Completed, you can proceed to step 2.

9. In step 2:

- Click on Files.
- On the Setup Connectivity Client page, select the operating system of your cluster.
- Download the files mentioned under Files.
- Follow the installation steps described under How to Install. You will need to SSH to the host mentioned in these instructions and run the commands.
- Once done, click Close.
- Click on Test Connection.
- Once the connection is successful, step 2 is marked as Completed. You can proceed to step 3.

10. In step 3:

- Click on Switch.
- A pop-up window appears with instructions on how to clean up CCMv1 resources. Under Select OS, select your operating system and then follow the cleanup steps described for your cluster type. You need to SSH to the specific cluster node indicated in the instructions in order to perform the cleanup. The steps that you need to perform are as follows (The same steps are printed in the UI):

- Centos7/RHEL7

On a CDH or CDP Private Cloud Base cluster, run the following on the Cloudera Manager host:

```
systemctl stop ccm-tunnel@CM.service
yum remove autossh
```

On an HDP cluster, run the following on the KNOX host:

```
systemctl stop ccm-tunnel@KNOX.service
yum remove autossh
```

- Centos6/RHEL6


On a CDH or CDP Private Cloud Base cluster, run the following on the Cloudera Manager host:

```
service reverse-tunnel stop 'CM'
yum remove autossh
```






On an HDP cluster, run the following on the KNOX host:

```
service reverse-tunnel stop KNOX
yum remove autossh
```

11.

Once your cluster has been upgraded the  icon appears next to its name.

12. From the ICON menu of the upgraded cluster select Refresh to synchronize the cluster with CDP:

Type	Status	Name	IP Address	Data Center	Port	Last Updated ↓	
	Connecting	Cluster 1 	172.27.92.129	CDH-CCMV1	7183	10/19/21 3:40 PM	⋮
	Active	Cluster 1 	172.27.98.65	CDH-6	7180	10/19/21 3:27 PM	Launch Cluster Manager
	Active	cd1 	172.27.69.199	HDP-CCMV1	8443	10/19/21 3:13 PM	Refresh Edit Remove Manage Access

13. The cluster status changes to Active.**Results**

Once you have performed the upgrade, you can continue using your cluster as usual.