

New UI for adding a CDP Private Cloud Base cluster for use in Replication Manager and Data Catalog (Preview)

Date published: 2022-03-29

Date modified: 2022-03-29

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Legal Notice

© Cloudera Inc. 2022. 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.

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.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Summary

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 services.

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. 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.

Note: A multi-cluster scenario is not supported by the Data Catalog use case. It is only supported by the Replication Manager use case.

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”.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided ‘as is’ without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

- 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.
5. 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.
6. On the Classic Clusters page, click Files in the Step 2 pane.
7. Follow the instructions in the *Setup Connectivity Client* dialog box. You need to download the `jumpgate-agentrpm` 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.

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

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

12. 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.
- e. In the **Knox Topology Configuration** section, you have two options: automatic deployment or manual deployment.

If you choose automatic deployment, simply provide Knox username and password and the token topology and gateway topology will be added automatically after clicking the Add button.

If you prefer not to provide Knox credentials, choose manual deployment and following the steps below:

- i. Click **Add** to proceed and on the next page.
- ii. Click on **Download Topology File** and a pop-up window appears.
- iii. Download the cdp-token file onto your computer.
- iv. Copy the downloaded file to the /var/lib/knox/gateway/conf/topologies folder on the Knox host of the CDP Private Cloud Base cluster that you are registering.
- v. Manually create the cdp_default.xml gateway topology file with required services. The cdp_default.xml file should include CM-API, CM-UI, ATLAS, ATLAS-API, RANGERUI, and RANGER services. For instructions, see [Configuring Apache Knox Gateway UI](#).

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.

- vi. Once done click **Close**.
- vii. Click **Test Connection** to verify the connection.
- viii. If the connection was successful, the registration status changes to *Completed*.

13. Once the cluster has been successfully registered, its status on the summary page changes to *Active*. If you click on cluster details, its connectivity status is *Reachable*.

Result

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

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.