

CDP Private Cloud Data Services 1.5.0

CDP Private Cloud Experiences Environments

Date published: 2020-12-16

Date modified: 2023-01-25

CLUSTERA

<https://docs.cloudera.com/>

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

Environments.....	4
CDP Private Cloud environment prerequisites.....	4
Registering a CDP Private Cloud environment.....	4
Accessing a CDP Private Cloud environment.....	6
Managing the compute cluster.....	6
Updating the Kubernetes configuration.....	6
Downloading the Kubernetes Configuration (OpenShift deployments only).....	6
Deleting a CDP Private Cloud environment.....	7

Environments

In CDP, a private cloud environment is a logical entity that represents the association of your Private Cloud user account with multiple compute resources using which you can provision and manage workloads such as Data Warehouse, Data Engineering, and Machine Learning. You can register as many environments as you require.

Registering an environment provides CDP with access to your user account and identifies the resources in your user account that CDP services can access or provision. For private cloud environments, resources include compute clusters such as Kubernetes as well as Data Lake clusters in CDP. Compute workloads are deployed within these environments.

A workload receives access to a Kubernetes cluster for compute purposes and a Data Lake cluster for storage, metadata, and security purposes within the environment in which it is deployed. Administrators can define user permissions and set resource quotas in each environment.

CDP Private Cloud environment prerequisites

CDP Private Cloud requires that you set up specific account criteria prior to registering an environment.

Set up the following before registering an environment:

- The trust store file that is configured in Cloudera Manager must meet the following requirements:
 - The file must not be empty
 - The file must be in JKS format (not PKCS12)
 - The file must contain the Certificate Authorities that have signed the following certificates:
 - The certificates of all Cloudera Manager servers that the user plans to use as base clusters in CDP Private Cloud
 - The certificates of all external Postgres servers used by Hive Metastore services in all base clusters

You can access the trust store file in Cloudera Manager at Administration > Settings > Security > "Cloudera Manager TLS/SSL Client Trust Store File".

- Your base cluster must meet the following requirements:
 - Ranger configured and running
 - Atlas configured and running
 - HMS configured and running
 - HDFS available and warehouse root configured
- A Kubeconfig file with cluster admin privileges.
- The associated Cloudera Manager user name and password.
- The Apache Ozone service must be enabled on your CDP Private Cloud Base cluster that you are using for the Cloudera Data Engineering service.

Registering a CDP Private Cloud environment

After you have set up the CDP Private Cloud requirements, you can register the environment.

Steps

Management Console UI

1. Sign into the CDP console.
2. Click Environments.
3. On the Environments page, click Register Environment.

- On the Register Environment page, provide the required information.

Environment

Environments in CDP Private Cloud provide shared data, security, and governance (metadata) for your machine learning and data warehouse applications.

Property	Description
Environment Name	Enter a name for your environment. This name will be used to refer to this environment in CDP. Note: Cloudera Data Warehouse (CDW) service requires that you specify the environment name 45 characters long. This is because CDW uses a deterministic namespace and adds a prefix to the environment name. The length of the namespace ID after CDW applies a prefix to the Environment name, including the hyphen (-), should not exceed 63 characters.

Compute Cluster Resources

To run workloads, you must specify a Kubeconfig file to register a Kubernetes cluster with CDP Private Cloud.

Property	Description
Kubernetes Configurations	Click Upload Files, then select a Kubeconfig file to enable CDP to access a Kubernetes cluster.
Storage class	The storage class on the OpenShift cluster. If you do not specify this value, the default storage class is used.
Domain	The default domain suffix for workload applications.

Data Lake

A Data Lake refers to the shared security and governance services in a CDP Data Center cluster linked to a CDP Private Cloud environment, and managed by Cloudera Manager. To register an environment, CDP Private Cloud needs to access Cloudera Manager and its Data Lake services.

Parameter	Description
Cloudera Manager URL	The Cloudera Manager URL.
Cloudera Manager Admin Username	The Cloudera Manager administrator user name.
Cloudera Manager Admin Password	The Cloudera Manager administrator password.

- Under Data Lake, click Connect.

When CDP Private Cloud has successfully connected to Cloudera Manager, a confirmation message appears, along with the Data Lake cluster services.

- Click Register.

The environment page appears. The new environment is also listed on the Environments page.



Note: You can repeat steps 2 through 5 to register more environments. If required, you can select a different Data Lake cluster while registering the additional environment. However, make sure that the Cloudera Manager instance is the same that you had specified to create the initial environment.

CLI

You can use the following command to create a new environment:

```
cdp environments create-private-environment \
--environment-name <value> \
--address <value> \
--authentication-token <value> \
----cluster-names <value>
```

For a detailed description of the command properties, use `cdp environments create-private-environment --help`

Accessing a CDP Private Cloud environment

Once an environment exists, you can access it from the Management Console.

About this task

From the details available on the Environments page, you can access the Data Lake cluster, the Data Hub clusters running within the environment, and the Summary page listing the information.

Procedure

1. To access an existing environment, navigate to Management Console > Environments and click on your environment.
2. On the environment details page, you can access the shared Data Lake cluster services and Cloudera Manager.

Managing the compute cluster

Compute Cluster provides information about the Kubernetes version, number of nodes in the cluster, and cluster registration time. You can access the RedHat OpenShift dashboard from here. You can download and update the Kubernetes configuration file.

Updating the Kubernetes configuration

You must update your [kubeconfig](#) file before it expires to ensure continued access to the cluster. The Kubeconfig file contains the cluster access information and authentication information for the admin user.

Procedure

1. Navigate to the Management Console > Environments.
2. Click the title of your environment.
3. Click the Compute Cluster tab.
4. Click Actions > Update Kubernetes Configuration.
5. Upload the new Kubeconfig file using the Choose File option and click Save.

What to do next

You can see the updated date and time in the Registered section.

Downloading the Kubernetes Configuration (OpenShift deployments only)

On OpenShift deployments, you can download and view the [kubeconfig](#) file used to register your environment.

Procedure

1. Navigate to the Management Console > Environments.
2. Click the title of your environment.
3. Click the Compute Cluster tab.
4. Click Actions > Download Kubernetes Configuration.

What to do next

The Kubernetes configuration file is downloaded to your system.

Deleting a CDP Private Cloud environment

Deleting an environment also removes all the resources associated with that environment.

Steps

Management Console UI

1. On the Environments page, click the environment that you want to remove.
2. Select Actions > Delete Environment.
3. Click Delete to confirm the deletion of the selected environment.

CLI

You can use the following command to delete an environment:

```
cdp environments delete-environment \  
--environment-name <value> \  

```

For a detailed description of the command properties, use `cdp environments delete-environment --help`