

Cloudera Data Warehouse Private Cloud Environments

Date published: 2020-08-17

Date modified: 2022-09-16



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

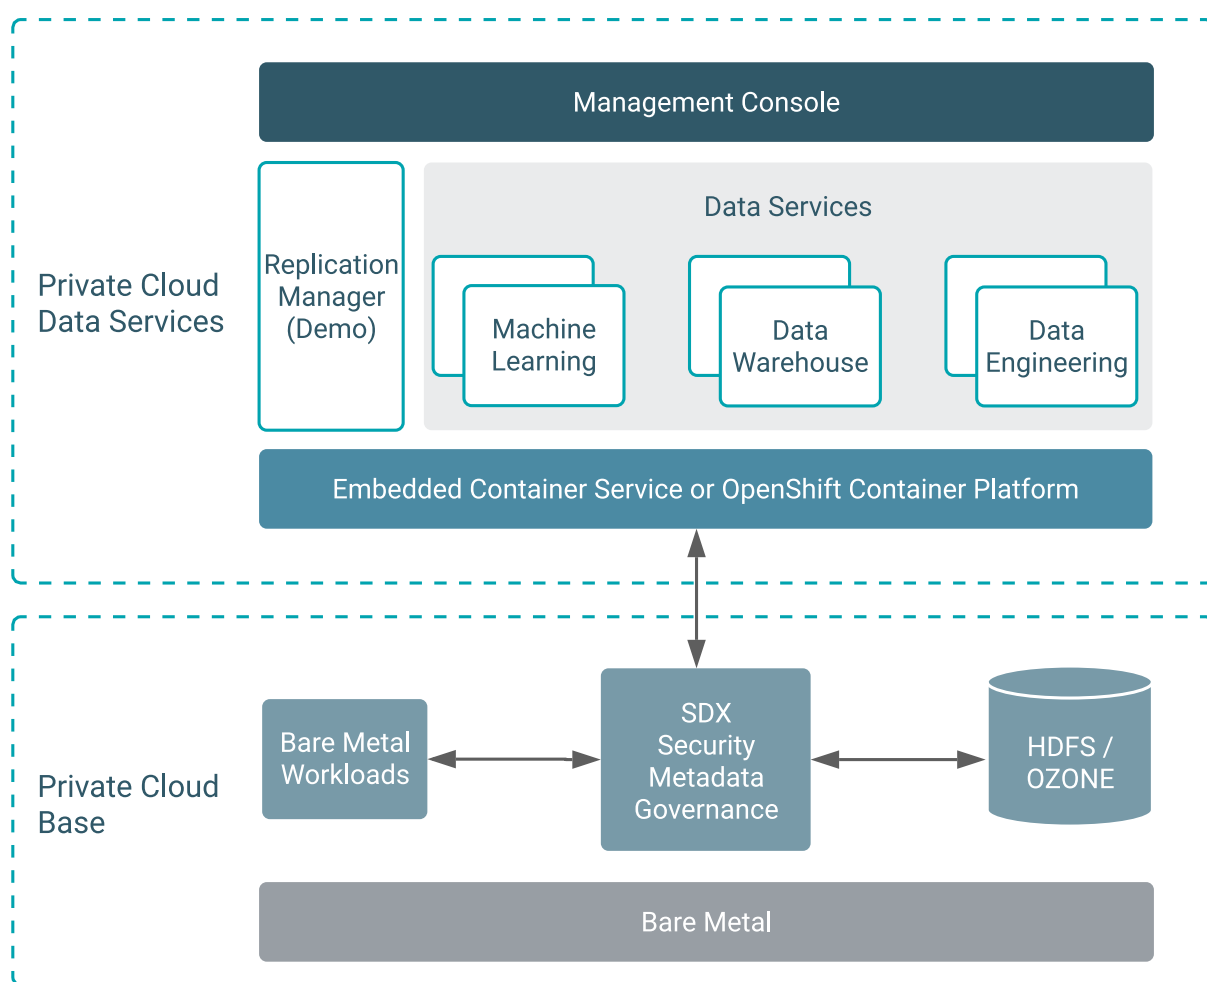
Overview of Private Cloud environments.....	4
Activating OpenShift and Embedded Container Service environments.....	5
Authentication using Delegation Username and Password.....	6
Viewing environment details in Cloudera Data Warehouse Private Cloud.....	7
Predefined Kerberos principals in Cloudera Data Warehouse Private Cloud.....	8

Overview of Private Cloud environments

An environment is a logical entity that represents the association of your Private Cloud user account with compute resources. Using the compute resources of the environment, you can provision and manage Cloudera Data Warehouse (CDW), Cloudera Data Engineering (CDE), or Cloudera Machine Learning (CML) workloads.

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

The following diagram shows the components of CDP Private Cloud:



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

Activating OpenShift and Embedded Container Service environments

This topic describes how to activate an environment to use for Cloudera Data Warehouse (CDW) Private Cloud.

About this task

Before you can create a Database Catalog to use with a Virtual Warehouse, you must activate a CDP environment. Activating an environment causes CDP to connect to the Kubernetes cluster, which provides the computing resources for the Database Catalog. In addition, activating an environment enables the Cloudera Data Warehouse (CDW) service to use the existing data lake that was set up for the environment, including all data, metadata, and security.

Before you begin

- Determine which environment that uses a particular data lake is the environment you want to activate for use with a Database Catalog and Virtual Warehouse.
- For local caching, ensure that an administrator uses the Local Storage Operator to create a local file system on an SSD/NVMe for each OpenShift worker node and then mounts it to a known location on the worker node. Make sure that this local caching location allows temporary data to be stored in a way that supports performance. You need to specify the Storage Class Name from the Local Storage Operator when you activate the environment for the CDW service in Step 4 below. For more information about creating a local file system on OpenShift worker nodes using the Local Storage Operator, see [Persistent storage using local volumes](#) in the OpenShift documentation.

In Embedded Container Service (ECS) environments, the Storage Class Name is automatically obtained from Cloudera Manager.

Procedure

1. In the CDW service, expand the Environments column by clicking the More... menu on the left side of the page.
2. In the Environments column, locate the environment that you want to activate.
3. When you locate the environment, click the activation icon to launch the Activation Settings dialog box.

- (For OpenShift environments) In the Activation Settings dialog box, specify the Storage Class Name from Local Storage Operator:

Activation Settings
✕

Do you want to activate the environment "dummy-openshift-env"?

Storage Class Name from Local Storage Operator *

Not a valid name

Security Context Constraint Name (optional)

Delegation Username* ⓘ

Delegation Password*

Pre-created database names for default database catalog* ⓘ

☐ Enable Low Resource Mode

CANCEL

ACTIVATE

This is the Storage Class Name you specified when you created the local file system for caching as described in the above section [Before you begin](#). It is the location where temporary data is stored.



Important: Be sure to specify the correct Storage Class Name when activating an environment. If an incorrect Storage Class Name is specified, the environment might activate successfully, but Virtual Warehouses that use the environment do not start.

Optionally, you can specify the Security Context Constraint Name.

- (For ECS environments) In the **Activation Settings** dialog box, enable low resource mode if required. You do not need to specify the Storage Class name or the Security Context Constraint name.
- Specify Delegation Username and Delegation Password to impersonate authorization requests from Hue to the Impala engine.
- If you are using an external database on your base cluster and want to use a default Database Catalog, then you must specify custom database name for DAS and Hue in the Pre-created database names for default database catalog field.
- Click **ACTIVATE**.

Authentication using Delegation Username and Password

You must specify a delegation username and password to impersonate authorization requests from Hue to the Impala engine during environment activation. The delegation user and password can authenticate users through an LDAP service account.

The ability to specify an LDAP delegation user also allows you to freely use special characters in your LDAP Bind DN, as CDW no longer has to inherit and process the delegation user from the LDAP Bind DN.



Note: You cannot update the delegation username and password after the environment is created. You must recreate the CDW environment to specify the delegation user.

The following image shows the CDW **Activation Settings** page containing the Delegation Username and Delegation Password fields:

Activation Settings ✕

Do you want to activate the environment "dummy-openshift-env"?

Storage Class Name from Local Storage Operator *

Not a valid name

Security Context Constraint Name (optional)

Delegation Username* ⓘ

Delegation Password*

Pre-created database names for default database catalog* ⓘ

☐ Enable Low Resource Mode

CANCEL

ACTIVATE

Viewing environment details in Cloudera Data Warehouse Private Cloud

This topic describes how to view details of an environment created on an OpenShift cluster in Cloudera Data Warehouse (CDW) Private Cloud.

About this task

You can view environment details without leaving the CDW Private Cloud service UI. Accessing the Environment Details page also enables you to edit the description of the environment.

Before you begin

You must activate an environment before you can view its details. See "Activating environments," which is linked to in the "Related information" section at the bottom of this page.

Procedure

1. In the CDW Data Service, expand the Environments column by clicking the More... menu on the left side of the page.
2. In the Environments column, click the search icon and locate the environment you want to view.

3. When you locate the environment you want to view, click the edit icon in the upper right corner of the tile. This launches the Environment Details page.
4. In the Environment Details page, you can view information about the environment, like when it was created and last updated, or how many Database Catalogs and Virtual Warehouses use the environment.
5. If you changed the environment description, click Apply in the upper right corner of the page to save it.

Related Information

[Activating OpenShift and Embedded Container Service environments](#)

Predefined Kerberos principals in Cloudera Data Warehouse Private Cloud

By default, Cloudera Data Warehouse (CDW) creates Kerberos principal names using the pod hostname when you create an environment. However, you can generate and provide the keytabs, if needed.



Attention: This feature is in technical preview and not recommended for use in production environments.

The service principals for CDW need to be the same as on the base cluster. For more information, see Customizing Kerberos principals in the CDP Private Base documentation.

By default, the host principals are generated programmatically. You can generate and provide the keytabs, but the hostnames in the Kerberos principals are fixed. For assistance, contact Cloudera Support.

Related Information

[Customizing Kerberos principals](#)