Cloudera Data Warehouse Private Cloud 1.4.0

Data Warehouse Release Notes

Date published: 2020-08-17 Date modified: 2022-09-16



Legal Notice

© Cloudera Inc. 2025. All rights reserved.

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

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

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

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

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

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

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

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

Contents

What's new in Cloudera Data Warehouse Private Cloud?	4
September 16, 2022	
June 22, 2022	
April 13, 2022	
January 14, 2022	
October 29, 2021	
October 4, 2021	
April 28, 2021	8
December 16, 2020	
August 17, 2020	
Known issues and limitations in Cloudera Data Warehouse Private Cloud	
General known issues	
Known issues on OpenShift cluster environments	
ECS cluster environments	
Known issues in Database Catalogs	
Known issues in Hive Virtual Warehouses	
Known issues in Impala Virtual Warehouses	
Known issues in Hue	
Data Analytics Studio	18
Fixed issues in Cloudera Data Warehouse Private Cloud (1.4.0)	18
Fixed issues in Cloudera Data Warehouse Private Cloud (1.4.0-H1)	19
Runtime component versions for Cloudera Data Warehouse Private Cloud	20

What's new in Cloudera Data Warehouse Private Cloud?

This section lists major features and updates for the Cloudera Data Warehouse (CDW) service in CDP Private Cloud.

September 16, 2022

There are no new features in the 1.4.0-H1 release of Cloudera Data Warehouse (CDW) service on CDP Private Cloud Data Services. For a list of fixes, see Fixed issues in Cloudera Data Warehouse Private Cloud (1.4.0-H1).

June 22, 2022

This release of Cloudera Data Warehouse (CDW) service on CDP Private Cloud Data Services introduces the new features and improvements that are described in this section.

Query isolation for scan-heavy, data-intensive queries in Hive LLAP Virtual Warehouses

Hive Virtual Warehouses base auto-scaling on the total scan size of the query. HiveServer, which receives all incoming queries, has a query planner component. When the HiveServer query planner receives queries, it examines the total scan size of each query. That is, it looks at the number of bytes read from the file system required to execute the query. If the Query Isolation feature has been enabled for a Virtual Warehouse and a query scans more data than the threshold set in the hive.query.isolation.scan.size.threshold parameter, the planner runs the query in isolation. This means that an isolated standalone executor group is spawned to run the data-intensive query. For more information, see Hive query isolation for ETL jobs.

Generate and download Hive diagnostic bundles

You can now generate and download diagnostic bundles containing log files for Hive Virtual Warehouses. For more information, see Downloading Hive diagnostic bundles in Data Warehouse Private Cloud.

SSL support added for MariaDB and MySQL databases

CDW can connect to SSL-enabled MariaDB and MySQL databases on the base cluster in addition to the SSL-enabled PostgreSQL database. For optimum security, the network connection between the default Database Catalog Hive MetaStore (HMS) in CDW and the relational database hosting the base cluster's HMS must be encrypted with SSL. See How to enable SSL support for MariaDB, MySQL, and Oracle databases.

SSL support added for Oracle databases (Preview)

CDW can connect to SSL-enabled Oracle database on the base cluster. See How to enable SSL support for MariaDB, MySQL, and Oracle databases.



Note: Cloudera recommends that you explore this feature in development or test environments. SSL-enabled Oracle database is not recommended for use in production deployments.

Ability to configure Impala coordinator and executor pod size (Preview)

You can optimize the performance of your Impala Virtual Warehouse and resources used in an environment based on your hardware configuration by customizing the amount of resources allocated to the Impala coordinators, executors, and catalog daemons. This helps you to better leverage intra-query parallelism and achieve powerful compute clusters with fewer nodes. For more information, see Creating custom pod configurations for Impala Virtual Warehouses.



Note: This is a preview feature and not recommended for use in production environments. Cloudera recommends that you explore this feature in test or development environments.

Node-level monitoring capabilities on ECS (Preview)

You can now monitor node-level metrics using Grafana. The metrics include CPU usage, memory usage, network usage, and disk IO. For more information, see Monitoring Data Warehouse service resources with Grafana dashboards.



Note: This is a preview feature and not recommended for use in production environments. Cloudera recommends that you explore this feature in test or development environments.

Auto-shutdown Impala coordinators

When you create a Virtual Warehouse, you can configure Impala coordinators to automatically shutdown during idle periods. You can also set a delay before the coordinator shuts down. For more information, see Configuring Impala coordinator shutdown.

Data Visualization integration in Cloudera Data Warehouse (Preview)

CDW integrates Data Visualization for building graphic representations of data, dashboards, and visual applications based on CDW data, or other data sources you connect to. Authorized users can explore data using graphics such as, pie charts and histograms and collaborate using dashboards. BI analysts who can access your environment can use these features. See Deploying Data Visualization in CDW.



Note: This is a preview feature and not recommended for use in production environments. Cloudera recommends that you explore this feature in test or development environments.

Impala Debug Web UIs are available in CDW Private Cloud

In CDW Private Cloud, you can now use the Impala debug web UIs, which map to equivalent debug web UIs in Cloudera Manager as follows:

CDW Debug Web UI	Cloudera Manager Equivalent
Impala Catalog Web UI	Catalog Server Web UI
Impala Coordinator Web UI	Impala Daemon Web UI
Impala StateStore Web UI	StateStore Web UI

For more information about this feature, see Using the Web UI to debug Impala Virtual Warehouses.

Hue supports Hive Hybrid Procedural SQL in CDW

You can run Hive Hybrid Procedural SQL (HPL/SQL) using the Hue query editor in CDW. For more information, see How to run a stored procedure from Hue in Cloudera Data Warehouse.

Improved read performance of ORC tables by Impala

Continuous improvements in Impala's read performance of ORC tables.

Security improvements

- CVE-2021-44228 (Apache Log4j 2 vulnerability) has been addressed in CDW on CDP Private Cloud 1.4.0 by upgrading Apache Log4j 2 to version 2.17.1.
- CDW Docker containers, applications, and modules now run as non-root users, including Hue.

April 13, 2022

This release of Cloudera Data Warehouse (CDW) service on CDP Private Cloud Data Services introduces the new features and improvements that are described in this section.

Support for SSL-enabled MariaDB and MySQL database (Preview)

CDW can connect to SSL-enabled MariaDB and MySQL databases on the base cluster. This feature is not ready for production deployment. Cloudera encourages you to explore this feature in a non-production environments and provide feedback on your experiences.



Note: You need to contact Cloudera to have this feature enabled.

Security improvements

CDW Docker containers, applications, and modules now run as non-root users. However, Hue containers still use root user.

January 14, 2022

This release of Cloudera Data Warehouse (CDW) service on CDP Private Cloud Data Services introduces the new features and improvements that are described in this section.

CDW supports MySQL database for HMS on the base cluster

The Hive Metastore (HMS) on the CDP private base cluster supports MySQL databases in addition to PostgreSQL and MariaDB. For more information on database requirements, see Base cluster database requirements for Cloudera Data Warehouse Private Cloud.

Users can access Hue using SSO

After logging in to the CDP Management Console, you can automatically log in to Hue using Single-Sign On (SSO). You do not need to enter your credentials to access Hue. To use this feature, you must first upgrade the CDW service to 1.3.3 (if you are on a previous version) and also create a new Virtual Warehouse. Upgrading an existing Virtual Warehouse to a newer version does not enable SSO authentication for Hue. The first user logging in to Hue will be a superuser.

CDW supports encrypted HDFS

CDW components such as Hive, Impala, and Hue can now read and write encrypted data stored on HDFS on the base cluster. If you are using Impala on an encrypted source data, then ensure that the data written to disk temporarily during processing is encrypted to keep the data confidential even during processing. For more information, see Configuring Impala Virtual Warehouses to spill encrypted data to HDFS in CDW Private Cloud.

Support for using Hive user-defined functions (UDF) (Technical Preview)

You can export user-defined functions (UDF) to a JAR file from a Hadoop- and Hive-compatible Java project and store the JAR file on HDFS. Using Hive commands, you can register the UDF based on the JAR file, and call the UDF from a Hive query using Hue in CDW. For more information, see Creating a user-defined function in Cloudera Data Warehouse.

Support for creating Impala tables in Kudu (Technical Preview)

You can configure an Impala Virtual Warehouse to connect to Kudu and create Impala tables in Kudu using the Hue query editor. For more information, see Configuring Impala Virtual Warehouses to create Impala tables in Kudu in Cloudera Data Warehouse Private Cloud.

Support for using custom generated Kerberos keytabs (Technical Preview)

By default, the principals and keytabs in CDW are generated programmatically. You can generate and provide the keytabs, but the hostnames in the Kerberos principals are fixed. To learn more about this, see Predefined Kerberos principals in Cloudera Data Warehouse Private Cloud.

CDW collects and provides diagnostic bundles for Impala

You can trigger a job to generate workload logs from the CDW web interface for Impala Virtual Warehouses. The diagnostic bundles are collated and stored under the tmp directory on HDFS. You can access and download the logs using the Hue service from the base cluster. For more information, see Generating and downloading diagnostic bundles.

Security improvements

- CVE-2021-44228 (Apache Log4j 2 vulnerability) has been addressed in CDW on CDP Private Cloud 1.3.3 by upgrading Apache Log4j 2 to version 2.16.
- CDP now sets stricter file permissions for CDW secrets.

October 29, 2021

There are no new features in this release of Cloudera Data Warehouse (CDW) service on CDP Private Cloud Data Services

October 4, 2021

This release of Cloudera Data Warehouse (CDW) service on CDP Private Cloud Data Services introduces the new features and improvements that are described in this topic.

Embedded Container Service (ECS) support for Cloudera Data Warehouse (CDW)

ECS service creates and manages an embedded Kubernetes deployment for the CDW Data Service in your data center. Cloudera Manager handles the installation and configuration of ECS clusters while you manage and monitor your ECS clusters. For more information on ECS, see ECS documentation.

Custom Security Context Constraint name in OpenShift environments

CDW provides an option to specify a custom Security Context Constraint (SCC) name while activating an OpenShift environment. This allows CDW to use 1000 as the UID for the Hive user. The custom SCC is not deleted when you deactivate or delete the CDW environment.

CDW supports Search and Bind users from multiple LDAP organizational units (OU)

Workload services such as Hue, Impala, and Hive now support LDAP Search Bind authentication with the following limitations, mostly related to filtering:

- You must specify at least two filter conditions, such as (&(uid={0})(objectClass=person)).
- You must specify at least two group search filter, such as (&(memberUid={0})(objectClass=posixGroup)(gidNumber=4001)).
- Hive CLI does not use the specified LDAP filter.

- Hue and Hive can use the User Search base for Group searches as well.
- Hue expects the groups to have 'cn' attribute as the group name attribute.
- Do not specify OU in the LDAP group search or user search. Remove "ou=xyz" from "ou=xyz,dc=cloudera,d c=com", so that it reads "dc=cloudera,dc=com".
- The group filter must be strict and must only return one group.
- Data Analytics Studio (DAS) can be configured only with Simple Bind authentication and uses the LDAP Bind DN as the Bind pattern.

Changes to supported databases for Hive Metastore on base cluster (Technical Preview)

The Hive Metastore (HMS) on the CDP Private Base cluster supports MariaDB databases in addition to PostgreSQL. This feature is in Technical Preview and is not ready for production deployment. For more information on database requirements, see Base cluster database requirements for Cloudera Data Warehouse Private Cloud.

Ability to configure database name for Hue and DAS

You can specify a custom database name for Hue, and DAS during environment activation when creating a default Database Catalog with external databases, or while creating a non-default Database Catalog. For more information, see Adding a new Database Catalog.



Note: If you are using an embedded database, then these fields do not apply.

DAS is disabled when HMS on the base cluster is configured with a non-PostgreSQL database

DAS supports only PostgreSQL databases. If you have configured your base cluster with a non-PostgreSQL database and if your Virtual Warehouse is linked to a default Database Catalog, then DAS is not installed on the cluster.

Ability to specify Delegation Username and Password

CDW no longer inherits the delegation user from the LDAP Bind DN. You must specify Delegation Username and Delegation Password on the Activation Settings page while creating a CDW environment. The delegation user and password can authenticate users through an LDAP service account. For more information, see Using delegation username and password.

April 28, 2021

This release of Cloudera Data Warehouse (CDW) service on CDP Private Cloud introduces the new features and improvements that are described in this topic.

New features

- Support for Red Hat OpenShift version 4.6
- Low resource hardware requirements mode for OpenShift environments: In this release, you can deploy CDW Private Cloud using fewer nodes, memory, and processors to get the functional benefits of CDW on Private Cloud. For details, see Get started with OpenShift low resource mode requirements.



Note: Cloudera recommends that you use the Low Resource Mode option for proof of concept (POC) purposes only. This feature is not recommended for production deployment.

- Impala Virtual Warehouses now support ACID reads on ORC tables: Impala Virtual Warehouse users can now run SQL queries on ORC ACID tables.
- Hive ACID compaction fixes: Multiple stability fixes have been made to increase the efficiency of Hive ACID compaction operations.

Behavioural changes

• New hardware requirements for standard resource mode: For details, see Get started with OpenShift standard resource mode requirements

- Custom service accounts can be used when installing CDP Private Cloud 1.2 with the following limitations and workarounds:
 - Limitation (DWX-6845): Demo data fails to load.
 - Limitation (DWX-6954): Only Hue is supported as the SQL editor for both Hive and Impala Virtual Warehouses.
 - Workaround (DWX-5941, DWX-7152): When you use custom service accounts during CDP Private Cloud 1.2 installation, you must add the custom service account users to the HDFS, HBase, and Hadoop SQL services using the Ranger Admin Web UI in Cloudera Manager:
 - 1. In Cloudera Manager, select RangerRanger Admin Web UI to launch the Service Manager page.
 - 2. In the Service Manager page, select the edit icon for HDFS to edit the service:
 - **a.** In the Add New Configurations region of the Edit Service page, add the custom HDFS service account user to the Value listed for the following parameters:

tag.download.auth.users

policy.download.auth.users

Make sure you separate user names with a comma.

For example:



Add the custom service user to both configuration parameters.

- b. Click Save to save your changes and return to the Service Manager page.
- 3. On the Service Manager page, click the edit icon for HBASE to edit the service:
 - **a.** In the Add New Configurations region of the Edit Service page, add the custom HBase service account user to the following configuration parameters:

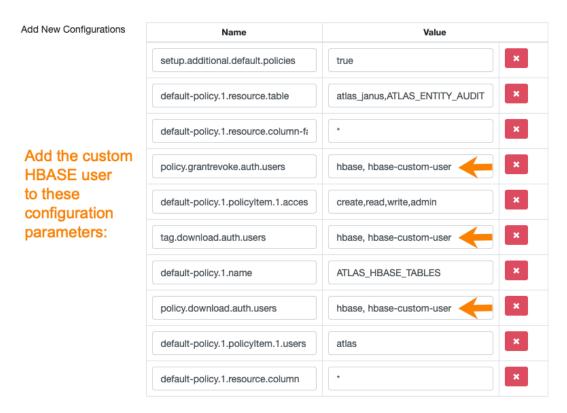
policy.grantrevoke.auth.users

tag.download.auth.users

policy.download.auth.users

Separate user names with a comma.

For example:



- **b.** Click Save to save your changes and return to the Service Manager page.
- **4.** In the Service Manager page, under HBASE, click the cm_hbase link to navigate to the configuration page for cm_hbase policies:



5. On the cm_hbase Policies page, click the edit icon for the all-table, column-family, column policy and add the custom HBase service account user to this policy:



- **6.** Click Save to save your changes and navigate back to the Service Manager page.
- 7. In the Service Manager page, click the edit icon for HADOOP SQL to edit the service:
 - a. In the Edit Service page, scroll down to the Add New Configurations region and add the custom Impala and Hive service account users, separating user names with commas, to the following configuration parameters:

tag.download.auth.users

policy.download.auth.users

policy.grantrevoke.auth.users

- **b.** To the policy.download.auth.users also add {USER}, separating user names with commas, to configure access to your SQL editor.
- c. Click Save to save your changes.

December 16, 2020

This release of Cloudera Data Warehouse (CDW) service on CDP Private Cloud introduces the new features and improvements that are described in this topic.

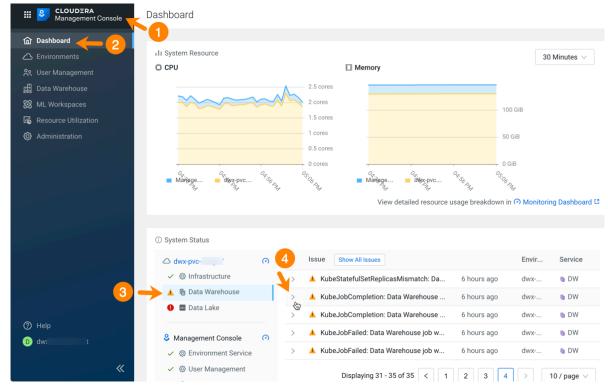
This release of CDW Private Cloud contains the following features:

- Support for Red Hat OpenShift version 4.5.
- Support for OpenLDAP in addition to Microsoft Active Directory LDAP.
- CDW endpoints on OpenShift clusters are now secured with TLS.



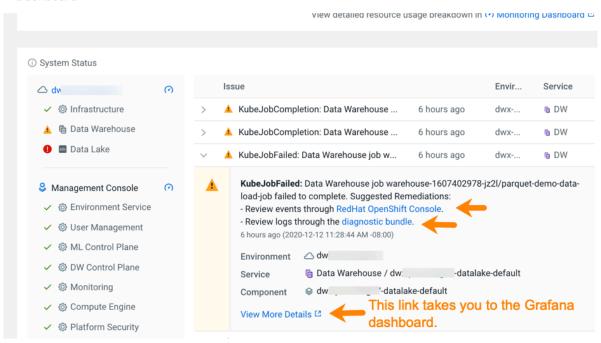
Tip: If you copied a URL to configure a connection in CDW Private Cloud 1.0, you must re-copy the new TLS-enabled URL in CDW Private Cloud 1.1. For example, you might have copied a URL to configure a connection in Data Analytics Studio, Hue, the JDBC URL, or a URL to connect with Impala shell. For information about managing Virtual Warehouse endpoints see, SSL-enabled client endpoints.

- Alerts have been added for the CDW service in the Management Console Dashboard to monitor the service health in terms of resource usage, network/security issues, and pod issues:
 - 1. Navigate to Management ConsoleDashboardData Warehouse, and click the arrow adjacent to the alert to expand it:



2. After expanding the alert, there are multiple options for troubleshooting. For example, in this KubeJobFailed alert for Data Warehouse, you can review the event in RedHat OpenShift Console, select the diagnostic bundle link to download logs, or click View More Details to launch the Grafana dashboard for this issue:

Dashboard



August 17, 2020

This release of Cloudera Data Warehouse (CDW) service on CDP Private Cloud introduces the new features and improvements that are described in this topic.

This is the first release of CDW service Private Cloud. It contains the following features:

- "Noisy neighbor" prevention achieved through workload isolation.
- Hive and Impala Virtual Warehouses with concurrency, auto-scaling, and auto-suspend features.
- Data Analytics Studio (DAS) and Hue SQL to support querying and analysis.
- On-premises LDAP integration and SSL with PostgreSQL.
- · Grafana monitoring dashboards for Hive and Impala.
- Diagnostics and log collection.
- Separate compute and storage on-premises.
- Improved CPU and memory utilization across the data warehousing apps.
- A central console from which you can secure and govern your data across data warehouses.

For an overview of CDW service Private Cloud and how to get started with CDW Private Cloud see Private cloud support and Getting Started with OpenShift requirements.

Known issues and limitations in Cloudera Data Warehouse Private Cloud

This section lists known issues and limitations that you might run into while using the Cloudera Data Warehouse (CDW) service in CDP Private Cloud Data Services.

General known issues

This topic describes the general platform-wide known issues for Cloudera Data Warehouse (CDW) Private Cloud. OPSAPS-58019: "SERVICE_PRINCIPAL is required for kinit" error while activating a new environment

If the /etc/krb5.conf file on the Cloudera Manager host contains "include" or "includedir" directives, then you may encounter Kerberos-related failures on both Embedded Container Service and Red Hat OpenShift platforms. You may see the following error in the Database Catalog's metastore pod logs: SERVICE_PRINCIPAL is required for kinit.

To resolve this issue:

- 1. SSH into the Cloudera Manager host as an administrator.
- 2. Open the /etc/krb5.conf file for editing.
- 3. Comment the lines containing "include" and "includedir" directives.
- 4. Save the changes and exit.
- 5. Recreate the environment.

DWX-10403: Executor pods get stuck in pending state with a warning

In rare circumstances, when Impala or Hive executors start up either due to autoscaling or by manually restarting the executors, the pods may get stuck in a pending state with a warning such as "volume node affinity conflict". This happens due to a race condition in the storage class that provides local volumes.

Restart the pods so that they can be rescheduled on new nodes with enough resources.

DWX-8525 and DWX-8526: CDW does not support upgrades from CDW 1.2 to 1.3.2

You may encounter issues while creating a new Virtual Warehouse on an existing Database Catalog (version 1.2) in an environment that is upgraded from 1.2 to 1.3.1.

The existing Virtual Warehouses will continue to operate. But if you want to create a new Virtual Warehouse, then you must reactivate the CDW environment after upgrading the base cluster from 1.2 to 1.3.2.

DWX-8502: HMS health check does not check port 9083

The HMS health check script does not check the health of its service port 9083 and may provide incorrect health status.

None.

DOCS-7431: LDAP limitations in CDW Private Cloud Virtual Warehouses

CDW Private Cloud workloads in Hive, Impala, Hue, and Data Analytics Studio (DAS) use simple bind authentication without filters. However, Management Console control plane uses search bind.

CDW Workloads use the LDAP Bind User field as the user bind pattern after replacing the user name with the pattern string. The workload users should be under the same sub-tree as the LDAP bind user.

DWX-5496: If Cloudera Manager truststore or the external database/vault is updated, CDW environment activation fails

Problem: If you have activated environments in CDW Private Cloud 1.1 and you update the Cloudera Manager truststore or the external database/vault truststore, you must update all of your CDW Private Cloud components else environment activation fails in CDW Private Cloud 1.1.

For example, you might perform truststore updates for the following use cases:

- To rotate the database certificate.
- To add a new CA certificate for a secure connection to an external vault.
- To rotate the Cloudera Manager certificate.

For information about these tasks, see Updating TLS certificates in the Management Console documentation set.

If any of these scenarios apply to your deployment, perform the following steps to make sure you can activate CDW environments for Database Catalogs and Virtual Warehouses in CDW Private Cloud 1.1:

- 1. In CDW Private Cloud 1.1, delete or remove all existing Virtual Warehouses.
- 2. Back up all the corresponding metadata or DDL to ensure you can recreate the objects contained in your Database Catalogs. Also back up your table data.



Important: Backup of corresponding metadata or DDL is mandatory for custom Database Catalogs, and it is also recommended for the default Database Catalog.

- 3. Delete or remove all existing Database Catalogs except the default Database Catalogs.
- 4. Deactivate the environment in CDW. To deactivate an environment, navigate to the environment

tile and click the deactivate icon $\stackrel{\searrow}{}$. Follow the prompts to deactivate the environment.

- 5. Update the certificate in the Management Console UI.
- **6.** In the CDW UI, reactivate the environment.
- 7. Recreate the non-default Database Catalogs.
- **8.** Recreate the Virtual Warehouses.

After you perform these steps, you can use your Virtual Warehouses with the updated truststore. For more information about managing SSL-enabled client endpoints in CDW Private Cloud, see SSL-enabled client endpoints.

Known issues on OpenShift cluster environments

This topic describes the OpenShift cluster environment known issues for Cloudera Data Warehouse Private Cloud.

DWX-4723: HTTP 504 Gateway timeout error returned when you delete an environment

When a CDW Private Cloud environment is deleted in the CDW Private Cloud UI, an HTTP 504 Gateway timeout error can be returned.

Using the OpenShift CLI, increase the route timeout to 60s on the OpenShift cluster:

```
oc annotate route <ROUTE-NAME> --overwrite haproxy.router.openshi ft.io/timeout=60s
```

For more information about setting this timeout, see the OpenShift documentation.

ECS cluster environments

This topic describes the Embedded Container Service (ECS) cluster environment known issues for Cloudera Data Warehouse Private Cloud.

BLESC-6074: Impala executor pods not running and queries from Hue fail

You may see the following error after submitting a query from Hue: Latest admission queue reason: Waiting for executors to start. Only DDL queries and queries scheduled only on the coordinator (either NUM_NODES set to 1 or when small query optimization is triggered) can currently run. You may also notice that the Impala executors are not running.

Make sure that the /etc/resolv.conf file on the ECS hosts contain a maximum of 2 search domains.

Known issues in Database Catalogs

This topic describes the Database Catalog known issues for Cloudera Data Warehouse Private Cloud.

DWX-14014 (Affects 1.4.0): Hue Query Processor port is hardcoded

The Hue Query Processor port is hardcoded to 5432, which does not allow the connection to an external PostgreSQL database using custom ports. Database Catalogs fail to start when you configure custom port for the external database during installation.

Use the default port (5432) or use the embedded database for the Database Catalog.

DWX-9245: Demo data set does not load on Containerized clusters

On Containerized clusters, the job to load the demo data may fail when you create a Database Catalog with the Load Demo Data option. As a result, you may not see the demo data set in Hue.

None.

DWX-8979: Parquet demo data set only loads partially on OpenShift clusters

On OpenShift clusters, the Parquet demo data set loads partially on a non-default Database Catalog, intermittently. As a result, you may intermittently and partially see the demo data set in Hue.

None.

Known issues in Hive Virtual Warehouses

This topic describes the Hive Virtual Warehouse known issues for Cloudera Data Warehouse (CDW) Private Cloud. **DWX-12613:** (Affected version: 1.4.0) Diagnostic bundle collection for Hive does not support custom time intervals and log types

You may see the options to select the time interval and the log types for which you want to generate the Hive diagnostic bundles, but the output may contain all the log files irrespective of your selection. This is a known issue, which would be fixed in a later release.

None.

DWX-8524: Hive queries on sys.db tables may fail

When you run queries on the sys.db table from Data Analytics Studio, you may see the following error: "Cannot create PoolableConnectionFactory (Could not open SSL root certificate file /m nt/config/current/hms-root.crt.)" on a default Database Catalog using PostgreSQL database. This happens because PostgreSQL requires SSL. This issue does not occur on Database Catalogs using MariaDB.

None.

DWX-6234: WAIT TIME seconds cannot be set to "0" in the Hive Virtual Warehouse UI

The slider that can be used to set WAIT TIME in the Hive Virtual Warehouse UI displays a range from 0 to 1000, but if you set it to 0, the UI automatically changes it to 60.

None.

DWX-4842: Entities are not being created in Atlas

Base clusters that are using Java 11 might be using truststores in PKCS12 format. Currently, Hive Virtual Warehouses on CDW Private Cloud only supports truststores in JKS format. This prevents the entities from being created in Atlas.

Using the keytool, convert the PKCS12 truststore in the base cluster to a JKS truststore.

Known issues in Impala Virtual Warehouses

This topic describes the Impala Virtual Warehouse known issues for Cloudera Data Warehouse Private Cloud. **DWX-13934:** (Affected version: 1.4.0) Impala returns 401 Unauthorized error when connecting to an Impala Virtual Warehouse

Hue and Impala shell display a 401 Unauthorized error when you submit a query to an Impala Virtual Warehouse. This issue can happen if the user who is trying to access Impala is not a part of an LDAP group.

Add the user to an LDAP group. Users must be part of at least one LDAP group for Impala.

DWX-10382: (Affected version: 1.4.0) Unable to connect to the pre-upgrade Impala Virtual Warehouse using impala-shell on an Embedded Container Service (ECS) cluster

After upgrading from 1.3.2, 1.3.3, or 1.3.4 to 1.4.0, you may not be able to access the pre-upgrade Impala coordinators through its endpoints.

Restarting the coordinator pods can solve the issue. Alternatively, you can patch the leader elector image in Impala coordinators after upgrading CDW as follows. This is necessary because the old leader elector image runs deprecated code which is not compatible with the latest RKE versions:

- 1. Start a terminal session that can connect to your Kubernetes cluster using kubeconfig, for example.
- 2. Run the following script:

```
#!/usr/bin/env bash
set -x

DWX_VERSION="1.4.0-b26"
NAMESPACE=${NAMESPACE:?please set the NAMESPACE to update}
LEADER_ELECTOR_IMAGE_PRE_UPGRADE=$(kubectl get sts coordina
tor -n $NAMESPACE \
   -o jsonpath={.spec.template.spec.containers[3].image})

LEADER_ELECTOR_IMAGE_POST_UPGRADE=$(sed \
   "s/:.*/:$DWX_VERSION/" <<< $LEADER_ELECTOR_IMAGE_PRE_UPGRA
DE)</pre>
```

```
kubectl patch sts coordinator -n $NAMESPACE --type='json' -
p='[{"op": "replace", "path": "/spec/template/spec/container
s/3/image", "value":"'$LEADER_ELECTOR_IMAGE_POST_UPGRADE'"}]'
kubectl delete pod -l app=coordinator
```

Known issues in Hue

This topic describes the Hue known issues for Cloudera Data Warehouse (CDW) Private Cloud.

DWX-12468 (Affected version: 1.4.0): Query store not reachable

You may see the following error when you click the **Queries** tab on Hue's **Job Browser** page. This is because the Hue Query Processor service is not yet available in CDW Private Cloud.

None.

DWX-9373: Unable to log into Hue as a local administrator

If you have logged into the CDP Management Console as a local administrator, then you may not be able to log into Hue and you may see the following error: "User is not authorized".

To access Hue, you must add your username to your organization's LDAP tree or log into the CDP Management Console using your LDAP credentials and then access Hue.

Data Analytics Studio

This topic describes the Data Analytics Studio (DAS) known issues for Cloudera Data Warehouse Private Cloud. **DWX-9034: Query search page does not display the query history**

The default value of the pod PID limit setting on Openshift is 1024. The das-event-processor pod in the Data Catalog's namespace may run out of the available processes over time on OpenShift, causing the query search page to not display query history. You may also see the following exception in the Event Processor logs: "ERROR com.hortonworks.hivestudio.eventProcessor.me ta.MetaInfoUpdater - Error occurred while reading dump metadata null".

You must manually restart the pod to resolve this issue. Alternatively, you can increase the pod PID limit or set it to maximum, which is a node level setting in Kubernetes/Openshift.

Fixed issues in Cloudera Data Warehouse Private Cloud (1.4.0)

This section lists issues fixed in this release of Cloudera Data Warehouse (CDW) service.

DWX-6132: Hive Virtual Warehouses require Cloudera Manager 7.2.4 and CDP 7.1.5

This issue has been fixed.

DWX-10126: CDW is unable to publish metadata messages to Kafka in a custom Kerberos-enabled cluster

This issue has been fixed.

DWX-8500: Toggling the auto-suspend option restarts all the pods on Hive Virtual Warehouses

This issue has been fixed.

DWX-10130: Unable to log out from Hue

This issue has been fixed.

Fixed issues in Cloudera Data Warehouse Private Cloud (1.4.0-H1)

This section lists issues fixed in this release of Cloudera Data Warehouse (CDW) service.

DWX-13017: Stats task during Hive query execution fails because of the missing "INSERTED_COUNT" column

Earlier, when Hive ran the stats task while a Hive query is running, Hive failed to update the transactional statistics. This issue has been resolved by incorporating a hotfix.

DWX-12801: Authorization error on SSO login in Cloudera Data Visualization on CDW Private Cloud

Cloudera Data Visualization displayed "Access Denied" error when you logged in using SSO. This issue can be resolved by creating a new Cloudera Data Visualization instance from CDW.

To create a new Cloudera Data Visualization instance:

- 1. Log in to Cloudera Data Warehouse as a DWAdmin.
- 2. Activate the environment.
- **3.** Create a Hive or an Impala Virtual Warehouse.
- 4. Create a Cloudera Data Visualization instance.
 - a. Click Data Visualization ADD NEW.
 - **b.** Specify a name for your instance, and select the environment and size from the drop-down menu and click CREATE.
- 5. After the Data Visulization instance is in the running state, open the application by clicking Open Data Visualization.

SSO is enabled on Cloudera Data Visualization in Private Cloud. If you are logged in to the Control Plane, then you will be logged into Cloudera Data Visualization automatically. If not, then log into the Control Plane using your LDAP credentials.



Note:

- Role Based Access Control (RBAC) is not yet available in Cloudera Data Visualization in Private Cloud. All users can access the Data Visualization instances as non-admin users by default.
- When you log in to Cloudera Data Visualization for the first time, you are added
 as a normal user. With these privileges, you can create datasets and visuals, but
 you cannot create new connections. To create new connections, you must have the
 "System Admin" role. By default, the admin user has the "Database admin" role.
- In CDP Private Cloud, the only available authentication mechanism is SAML SSO. Once you log out from Data Visualization, you cannot sign back into the application using your LDAP credentials. You must go to CDW and launch Cloudera Data Visualization.

To log in as an administrator and change your role:

1. Log out from the Cloudera Data Visualization application.

2. Log in to Cloudera Data Visualization using the following credentials:

Username: vizapps_admin Password: vizapps_admin



Note:

- vizapps_admin is the only available admin user. Cloudera recommends that you change the password for the vizapps_admin user.
- New users logging in to Cloudera Data Visualization are automatically added to the "viz_guest_group" group. You can assign the "System Admin" role to this group to grant new users the permision to create new connections.
- 3. On the Cloudera Data Visualization application, click Manage Users & Groups .
- 4. Click to assign the "System Admin" role to the required LDAP users.

You can also add new users or groups and assign roles to them.

To create new connections:

- 1. Launch Cloudera Data Visualization from CDW using LDAP credentials.
- 2. Go to the DATA tab and click NEW CONNECTION.

The **Create new Data Connection** dialog box is displayed.

3. Select the connection type, specify a connection name, and select the warehouse you want to connect to from the drop-down menu.

The corresponding connection details (hostname and port) are automatically populated.

Under Credentials, specify the username that you used to log in to Cloudera Data Visualization, as other users may not have proxy privileges in the warehouse, and click TEST.



Note: You need not enter the password because you are authenticated with SSO. Trusted impersonation is enabled in CDW.

If the connection test is successful, the "Connection Verified" message is displayed.



Note: (For connections to Impala Warehouse) If the connection test fails, retest the connection using the delegation username and password that you used while activating the CDW environment.

4. Click CONNECT.

The newly created connection is displayed.

Runtime component versions for Cloudera Data Warehouse Private Cloud

Cloudera Data Warehouse (CDW) uses Hive, Impala, and Hue as its Runtime components. The following table lists the versions of Hive, Impala, and Hue used in the CDW Private Cloud 1.4.0 release:

Runtime component	Version
Hive	2022.0.7.0-80
Impala	2022.0.7.0-80
Hue	2022.0.7.0-80