

Cloudera Data Warehouse Private Cloud 1.5.0

Data Warehouse Release Notes

Date published: 2020-08-17

Date modified: 2023-01-25

CLOUDERA

<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

What's new in Cloudera Data Warehouse Private Cloud.....	4
Known issues and limitations in Cloudera Data Warehouse Private	
Cloud.....	5
General known issues.....	5
Upgrade-related known issues.....	6
Known issues on OpenShift cluster environments.....	9
ECS cluster environments.....	9
Known issues in Database Catalogs.....	9
Known issues in Hive Virtual Warehouses.....	10
Known issues in Impala Virtual Warehouses.....	10
Known issues in Hue.....	11
Data Analytics Studio.....	12
Unified Analytics.....	12
Fixed issues in Cloudera Data Warehouse Private Cloud.....	13
Runtime component versions for Cloudera Data Warehouse Private	
Cloud.....	14

What's new in Cloudera Data Warehouse Private Cloud

Learn about the new features in Cloudera Data Warehouse (CDW) service on CDP Private Cloud Data Services 1.5.0.

Changes to the database recommendations

In older releases, Cloudera recommended that you use an external database for the Hive MetaStore (HMS) and the Control Plane service, so that you could backup and restore data as needed. With the introduction of the Data Recovery Service in 1.5.0, you no longer need to manually back up and restore the data in the database. Starting from 1.5.0, Cloudera recommends that you use an embedded database for the HMS and the Control Plane service.

CDW in Private Cloud supports Unified Analytics

Unified Analytics bring SQL equivalency without syntax changes to CDW SQL engines. Unified Analytics provide significant optimization equivalency to these engines, unifying common techniques such as subquery processing, join ordering, and materialized views. Unified Analytics documentation includes features, limitations, and how to use Unified Analytics. To take advantage of this feature, create a new Impala Virtual Warehouse and enable the Unified Analytics option, or simply create a new Hive Virtual Warehouse. For more information, see [Unified Analytics overview](#).

The Unified Analytics feature has the following limitations:

- Query isolation feature for Unified Analytics is in technical preview in the Private Cloud 1.5.0 release and is not recommended for use in production deployments. Cloudera recommends that you try this feature in test or development environments.
- Queries with left or right anti joins are not supported.

CDW in Private Cloud supports Apache Iceberg on HDFS (Preview)

Apache Iceberg integration with Cloudera Data Platform (CDP) enhances the Lakehouse architecture by extending multifunction analytics to a petabyte scale for multi-cloud and hybrid use cases. From Hive or Impala, you use Apache Iceberg features in CDW, which include time travel, create table as select (CTAS), and schema and partition evolution. To use Iceberg with CDW, you must upgrade to CDP Private Cloud Data Services 1.5.0.

Apache Iceberg V1 is in technical preview in the Private Cloud 1.5.0 release and is not recommended for production deployments. Cloudera recommends that you try this feature in test or development environments.



Note: Iceberg tables cannot be accessed from Hive, Impala, Spark, and Flink on the base cluster. You need to be on CDW.

For the list of supported features and more information, see [Iceberg overview](#).

Ability to back up and restore Kubernetes data using DRS

The Data Recovery Service (DRS) in CDP enables you to back up and restore Kubernetes namespaces behind CDW entities such as Database Catalogs and Virtual Warehouses on demand. CDW leverages DRS and provides CDP CLI endpoints which you can use to create and restore backups for CDW namespaces to back up CDW metadata and configurations such as Kubernetes objects, persistent volumes, autoscaling configuration, and so on. To learn more about the Data Recovery Service, see [Data Recovery Service overview](#). For the list of available sub-commands for CDW, see [Using DRS with CDW](#).

Support added for ADLS Gen2 object storage (Preview)

CDW supports using ADLS (Gen1 and Gen2) containers for storing tables. CDW exposes Hive and Impala tables stored on ABFS containers as SQL tables which you can query using Hue. However, you cannot browse and import files to create tables from ABFS in Hue. For more information, see [Supported object storage services for Cloudera Data Warehouse Private Cloud](#).

This feature is in Technical Preview and is not recommended for production deployments. Cloudera recommends that you try this feature in test or development environments.

Non-default Database Catalogs are deprecated

The ability to create non-default Database Catalogs has been deprecated and is disabled by default. Cloudera encourages you to use the default Database Catalog that is created when you activate an environment. However, if you need to use a non-default Database Catalog, then you can enable the Create multiple Database Catalogs option from the **Advanced Settings** page as described in [Enabling the option to create additional Database Catalogs in CDW Private Cloud](#).

Ability to migrate Hive workloads from CDP Private Cloud Base to CDW Data Service on CDP Private Cloud

You can migrate Hive workloads from CDP Private Cloud Base to CDW Private Cloud to leverage the auto-scaling, workload optimization, isolation, data caching, and many other powerful capabilities that CDW offers. For more information, see [Migrating Hive workloads from CDP Private Cloud Base to CDW Private Cloud](#).

Improved troubleshooting experience for failed Database Catalogs and Virtual Warehouses

A new option called View Logs is displayed on Database Catalogs and Virtual Warehouses that have failed. It enables you to quickly view the logs and start investigating the issue. This is in addition to the diagnostic bundles that you can generate on a need basis.

CDW supports Cloudera Data Visualization (CDV) 7.0.5

To know what's new in CDV 7.0.5, see [What's new in Data Visualization 7.0.5](#).

Support for using Hive user-defined functions (UDF) is GA

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](#).

This feature is generally available (GA) starting 1.5.0, and you can use it in production environments.

Support for uploading auxillary JARs

CDW enables administrators to upload auxillary JARs to the Hive classpath that might be required to support dependency JARs, third-party Serde, or any Hive extensions. For more information, see [Uploading additional JARs to CDW](#)

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 service-wide known issues for Cloudera Data Warehouse (CDW) Private Cloud.

Known Issues identified in 1.5.0

FreeIPA is not supported

Cloudera does not support FreeIPA as a certified identity management system for CDW data service on CDP Private Cloud 1.5.0.

Known Issues identified in 1.4.1

DWX-13759: DAS server and DAS WebApp pods exists even after disabling DAS

After you disable DAS from the **Advanced Settings** page, you may still see the DAS server and DAS WebApp pods running in the Database Catalog and Virtual Warehouse namespaces respectively.

Delete and recreate the Database Catalog and Virtual Warehouse.

Known Issues identified before 1.4.1

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-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.

Upgrade-related known issues

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

DWX-8525 and DWX-8526: CDW supports upgrades from previous releases to 1.5.0 release with listed limitations

You may encounter issues while creating a new Virtual Warehouse on an existing Database Catalog (for example version 1.4.1) in an environment that is upgraded from, say 1.4.1 to 1.5.0.

The existing Virtual Warehouses continue to operate. But to create a new Virtual Warehouse, you must reactivate the CDW environment after upgrading the Private Cloud data services from a previous release to a newer release (for example from 1.4.1 to 1.5.0.)



Important: Review the following caveats before upgrading to 1.5.0:

- Reactivating the environment deletes the existing Database Catalogs and Virtual Warehouses. However, the metadata stored in the Hive Metastore (HMS) and data that is stored on HDFS or Ozone from the previous version is not deleted from the base cluster. Cloudera recommends that you manually take a backup of the Virtual Warehouse configurations, which you can use to create new Virtual Warehouses after the upgrade.
- (For ECS) On ECS, a “default” environment is created by the Control Plane when you add a Private Cloud cluster. If you have more than one ECS clusters managed using different instances of Cloudera manager, but using the same Active Directory (AD) server and using the same “default” environment, then go to the [Advanced Configuration Advanced Settings](#) page and disable the Use deterministic namespace names option after performing the upgrade but before reactivating the environment in CDW.

If you did not use the default environment which is created in CDW, then you can create a new environment from Management Console with a unique name.

- (For OCP) If you have more than one OCP clusters managed using different instances of Cloudera Manager, but using the same AD server and using the same environment name, then go to the [Advanced Configuration Advanced Settings](#) page and disable the Use deterministic namespace names option after performing the upgrade but before reactivating the environment in CDW.

If you are creating a new environment from the Management Console with a unique name, then you need not disable the Use deterministic namespace names option.

- Apply workaround for the AD limitation (TSB-639): If Cloudera Manager does not manage the `/etc/kerb5.conf` file, then add the following lines in the `/etc/kerb5.conf` file manually to every host in the cluster:

```
[libdefaults]
allow_weak_crypto = true
```

DWX-14192: Atlas entities not created after upgrading from 1.4.0-H1 to 1.5.0

After upgrading from 1.4.0-H1 to 1.5.0, you may notice that Atlas entities are not created and you may see the following error in the metastore pod logs: `java.security.UnrecoverableKeyException: Password verification failed`

1. Log in to the Data Warehouse service as DWAdmin.
2. Go to the Database Catalog tile and note the Database Catalog namespace.
3. Get the Kubernetes configuration for accessing the cluster using the `kubectl` command.

For OCP, see [Downloading the Kubernetes Configuration](#).

On ECS deployments, log into the nodes as a root user. Login into the master node, which is usually the first node, and check access to the cluster with the following `kubectl` command:

```
/var/lib/rancher/rke2/bin/kubectl --kubeconfig=/etc/rancher/rke2/rke2.yaml get nodes
```



Note: On ECS, run the `/var/lib/rancher/rke2/bin/kubectl` command on the node you accessed in the SSH session, and use the `/etc/rancher/rke2/rke2.yaml` kubeconfig file in the next steps.

4. Run the following command to get the `jcks-secret`:

```
kubectl get secret jcks-secret -n [***DBC-NAMESPACE***] -o yml --kubeconfig=/path/to/kubeconfig
```

5. Create a backup of the jceks-secret file as follows:

```
kubectl get secret jceks-secret -n [***DBC-NAMESPACE***] -o
yaml --kubeconfig=/path/to/kubeconfig > jceks-secret.backup.
yaml
```

6. Verify that the contents of the backup file are same as the original file as follows:

```
cat ./jceks-secret.backup.yaml
```

7. Save the jceks-secret.backup.yaml file.
8. Create a new jceks-secret file with the updated password while specifying the “metastore-0” pod as follows:

```
kubectl exec -it metastore-0 -n [***DBC-NAMESPACE***] --kube
config=/path/to/kubeconfig -- /bin/bash -c "cp /jceks/secret
s.jceks /tmp/ && hadoop credential delete truststore.passwor
d -provider jceks://file/tmp/secrets.jceks -f && hadoop cred
ential create truststore.password -value changeit -provider
jceks://file/tmp/secrets.jceks && cat /tmp/secrets.jceks | b
ase64 -w 0 > /tmp/secrets.jceks.b64 && hadoop fs -rm file:///
tmp/secrets.jceks"
```

This creates a copy of the secrets.jceks file, removes the truststore.password key from the tmp/secrets.jceks file, creates a new /tmp/secrets.jceks.b64 file on the filesystem of the metastore-0 pod, and finally removes the /tmp/secrets.jceks file from the filesystem of the metastore-0 pod.

The /tmp/secrets.jceks.b64 file is used for creating a patch file in the next step.

9. Create a patch file as follows:

```
echo "data:" > jceks-secret.patch.yaml && echo -n " secrets
.jceks: " >> jceks-secret.patch.yaml
```

The jceks-secret.patch.yaml patch file is required for modifying the jceks secret that is stored in the Kubernetes environment.

10. Run the following command to append data from the /tmp/secrets.jceks.b64 file (present on the metastore-0 pod) to the jceks-secret.patch.yaml file (present on the local filesystem):

```
kubectl exec -it metastore-0 -n [***DBC-NAMESPACE***] --kube
config=/path/to/kubeconfig -- /bin/bash -c "cat /tmp/secrets
.jceks.b64" >> jceks-secret.patch.yaml
```

11. Verify the contents of the jceks-secret.patch.yaml file as follows:

```
cat jceks-secret.patch.yaml
```

12. Apply the patch as follows:

```
kubectl patch secret jceks-secret -n [***DBC-NAMESPACE***] -
-kubeconfig=/path/to/kubeconfig --patch-file ./jceks-secret.
patch.yaml
```

13. Restart the metastore pods as follows:

```
kubectl delete pod metastore-0 -n [***DBC-NAMESPACE***] --ku
beconfig=/path/to/kubeconfig
```

```
kubectl delete pod metastore-1 -n [***DBC-NAMESPACE***] --ku
beconfig=/path/to/kubeconfig
```

Known issues on OpenShift cluster environments

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

Known Issues identified in 1.5.0

No new known issues in this area.

Known Issues identified before 1.4.1

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.openshift.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.

Known Issues identified in 1.5.0

No new known issues in this area.

Known Issues identified before 1.4.1

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.

Known Issues identified in 1.5.0

DWX-14281: Configuration changes to Database Catalogs are not applied

When you change Database Catalog configurations from the CDW UI, the changes are not applied.
None.

DWX-14022: The metastore-sys-init-job job fails with "Table not found 'TXNS'" error

You see that the default Database Catalog starts successfully and the UI indicates it as running, however, the metastore-sys-init pods are in an error state. You also see that the Hive MetaStore (HMS) pods fail with the following error: "[HiveServer2-Background-Pool: Thread-125] ERROR org.apache.hadoop.hive ql.parse.CalcitePlanner - org.apache.hadoop.hive ql.parse.SemanticException: Line 1:1117 Table not found 'TXNS'". This happens because the TXNS table is not present on CDP Private Cloud Base version 7.1.7 SP2.

You must upgrade the base cluster to CDP Private Cloud Base version 7.1.8 and then rebuild the affected Database Catalog.

Known Issues identified in 1.4.1

DWX-13758: Updating the pre-upgrade Database Catalog and Virtual Warehouse fails silently

If you update or refresh the Database Catalogs and Virtual Warehouses that were created before you upgraded to 1.4.1 by changing their configurations, then they may go into an erroneous state.

Do not update or refresh the Database Catalogs and Virtual Warehouses that were created before you upgraded to 1.4.1.

Create new Database Catalogs and Virtual Warehouses.

Known Issues identified before 1.4.1

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.

Known Issues identified in 1.5.0

DWX-14302: Hive query isolation executor pod crash loops with "Failed to start LLAP Daemon" exception

If you have enabled the Query Isolation option on a Hive Virtual Warehouse, the queries that should run in isolation may fail to start the executor, and eventually time out with an exception.

None.

Known Issues identified before 1.4.1

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.

Known Issues identified in 1.5.0

No new known issues in this area.

Known Issues in 1.4.0

DWX-13934: 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.

Known issues in Hue

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

Known issues identified in 1.5.0

DWX-12616: Hue limitation in CDW Private Cloud

Following are the known limitations in Hue in CDW Private Cloud 1.5.0:

- Hue Importer does not support importing files more than 200 KB in size
- Hue File Browser does not work if the HDFS service on the base cluster is configured for high availability

None.

DWX-13865: Hue File Browser does not work with HDFS HA

Hue File Browser is not accessible or displays a 403 error when you click on File Browser from the left assist panel in CDW Private Cloud if HDFS is configured for High Availability on the base cluster. Currently, Hue in CDW cannot obtain the hostname and the port from the HttpFS service on the base cluster. This is a known limitation.

You must manually build and specify the WebHDFS URL for Hue in CDW to connect to the HttpFS service on the base cluster.

1. Log in to Cloudera Manager as an Administrator.
2. Go to **Clusters Hive Configuration** and note the value present in the Kerberos Principal field.

This is the Hive service account name.

3. Go to **Clusters HDFS HttpFS Advanced Configuration Snippet (Safety Valve) for httpfs-site.xml** and click  to add the following lines:

Name: httpfs.proxyuser.hive.hosts, Value: *

Name: httpfs.proxyuser.hive.groups, Value: *

Replace hive with the actual Hive service account name.

4. Click **Save Changes** and restart the HDFS service or the HttpFS role.
5. Go to the **Instances** tab and note the hostname of the HttpFS role.
6. Go to the **Configuration** tab and note the port for the `hdfs.httpfs.http.port` configuration from the **RESTPort** field.

The default value of the `hdfs.httpfs.http.port` configuration is 14000.

7. Use the hostname and the port to construct the WebHDFS URL as follows:

```
https://[***HOSTNAME***]:[***PORT***]/webhdfs/v1
```

8. Log in to the Data Warehouse service as a DWAdmin.
9. Go to the Virtual Warehouse from which you want to connect Hue to the base cluster's HDFS service and click  **Edit**.
10. Go to **CONFIGURATIONS Hue**, select `hue-safety-valve` from the **Configuration files** drop-down list and add the following lines:

```
[hadoop]
```

```
[[hdfs_clusters]]
[[[default]]]
webhdfs_url=https://[***HOSTNAME***]:[***PORT***]/webhdfs/v1
```

Specify the WebHDFS URL that you constructed earlier.

11. Click Apply changes.

Update the hue-safety-valve configuration for any Hive or Impala Virtual Warehouses from which you want to connect to the base cluster HDFS.

Known Issues identified before 1.4.1

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.

Known Issues identified in 1.5.0

No new known issues in this area.

Known Issues identified before 1.4.1

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.meta.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.

Unified Analytics

This topic describes the Unified Analytics known issues for Cloudera Data Warehouse Private Cloud.

Known issues identified in 1.5.0

DWX-13362: Queries containing simple SELECT statements get cached

A query such as `SELECT * FROM planes LIMIT 2;` gets cached. This is an undesirable result and will be fixed in a future release.

None.

DWX-8348: Service Discovery does not work when Impala is enabled with Unified Analytics

Impala provides multiple endpoints for clients, such as Impala coordinator, impala-proxy, and Hive Server2 (HS2) (when Impala Virtual Warehouse is used in Unified Analytics mode). When Impala is enabled with Unified Analytics, both coordinator and HS2 endpoints are active, which interferes with Service Discovery.

None.

DWX-13891: Hive queries intermittently fail with “InvalidACLException” error

You may see the following error in Hue or Beeline when you run the DROP DATABASE, DROP TABLE, or ALTER TABLE DROP PARTITION operations on an Impala Virtual Warehouse running in the Unified Analytics mode: org.apache.zookeeper.KeeperException \$InvalidACLException: KeeperErrorCode = InvalidACL for /llap-sasl/user-hive.

To resolve this issue perform the following steps:

1. Log in to the Data Warehouse service as DWAdmin.
2. Go to Virtual Warehouse tile  Edit CONFIGURATIONS Hiveserver2 and select hive-site from the Configuration files drop-down menu.
3. Click .
4. Add the following line on the **Add Custom Configurations** pop-up and click Add:
`hive.llap.io.proactive.eviction.enabled=false`
5. Click Apply Changes.

Fixed issues in Cloudera Data Warehouse Private Cloud

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

DWX-13718: Unable to register UDFs in CDW using the auxiliary JARs approach

This issue has been fixed.

DWX-14014: Hue Query Processor port is hardcoded

The Hue Query Processor port is no longer hardcoded to 5432, which now allows you to connect to an external PostgreSQL database using custom ports.

DWX-13813: Unable to select or delete custom pod configurations for Impala

This issue has been fixed.

DWX-13816: Refresh option at the environment level is not functional

This issue has been fixed.

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

This issue has been fixed.

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

This issue has been fixed.

ENGESC-18302: Unable to collect diagnostic bundles in CDW

Earlier, there were issues pulling the diagnostic-data-generator image from the registry and deleting the diagnostic bundle collection jobs from the CDW web interface.

This issue has been fixed.

DWX-13818: Impala Virtual Warehouse shows incorrect executor count on the CDW UI when the executor is stuck in the pending state

This issue has been fixed.

DWX-13788: Hive Virtual Warehouse is stuck in stopped state and displays the hive server service is not ready warning when you refresh or update the Virtual Warehouse

This issue has been fixed.

DWX-13095: Hive Virtual Warehouse needs to be restarted after 7 days

If a Hive Virtual Warehouse ran continuously for 7 days, you may have seen that every Hive query fails with the following error: LLAP_TOKEN can't be found in cache.

This issue has been fixed.

DWX-13564: Remove cleartext passwords from CDW configurations

CDW configurations such as `hive.metastore.dbaccess.ssl.properties`, `hadoop.security.group.mapping.ldap.ssl.truststore.password`, `hadoop.security.group.mapping.ldap.bind.password`, `ssl.client.truststore.password`, and so on now obtain the passwords from the keystore secret and no longer have passwords in cleartext.

This makes CDW more secure because passwords are not visible or configurable from the CDW web interface.

DWX-13646: Unable to create Database Catalogs or Virtual Warehouses by specifying the image version using CDP CLI commands

Creating Database Catalogs or Virtual Warehouses using CDP CLI commands while specifying the image version is not supported in CDW on Private Cloud. The CDP CLI 0.9.76 release and the corresponding reference documentation have been updated to reflect this change. Also, you get the following message when you try this: The argument `.imageVersion` is not available under the private CDP form factor. It is only available for these form factors: `public`. Check that your profile configuration or explicit endpoint URL points to the correct control plane.

This issue has been fixed.

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 1.5.0 release:

Runtime component	Version
Hive	2022.0.11.1-15
Impala	2022.0.11.1-15
Hue	2022.0.11.1-15