

Cloudera Runtime 7.2.16

Using Hue

Date published: 2020-07-28

Date modified: 2023-01-11

CLOUDERA

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

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

About using Hue.....	4
Accessing and using Hue in Cloudera Data Warehouse.....	4
Viewing Hive query details.....	6
Viewing Hive query history.....	6
Viewing Hive query information.....	6
Viewing explain plan for a Hive query.....	7
Viewing Hive query timeline.....	8
Viewing configurations for a Hive query.....	8
Viewing DAG information for a Hive query.....	10
Terminating Hive queries.....	13
Comparing Hive queries in Hue.....	13
How to run a stored procedure from Hue in Cloudera Data Warehouse.....	14
Enabling stored procedures for Hive in Cloudera Data Warehouse.....	15
Enabling the SQL editor autocompleter.....	16
Using governance-based data discovery.....	16
Searching metadata tags.....	17
Using Amazon S3 with Hue.....	17
Using Azure Data Lake Storage Gen2 with Hue.....	18
Granting permission to access S3 and ABFS File Browser in Hue.....	18
Supported non-ASCII and special characters in Hue.....	18

About using Hue

Hue provides a one-stop querying experience in Cloudera Data Platform (CDP) that leverages Hive, Impala, and Unified Analytics SQL engines.

Accessing and using Hue in Cloudera Data Warehouse

Get started using Hue by analyzing and visualizing your data with Impala and Hive SQL query engines.

About this task

To try Hue without having an account, try running sample queries on <http://demo.gethue.com/>.

Before you begin

Hue uses your LDAP credentials that you have configured for the CDP cluster.

Procedure

1. Log into the CDP web interface and navigate to the Data Warehouse service.
2. In the Data Warehouse service, navigate to the **Overview** page.




Note: You can also launch Hue from the **Virtual Warehouse** page using the same steps.

3. To run Impala queries:

- a) On the **Overview** page under Virtual Warehouses, click on the Hue button.

The query editor is displayed:



- b) Click a database to view the tables it contains.

When you click a database, it sets it as the target of your query in the main query editor panel.

- c) Type a query in the editor panel and click the run icon ▶ to run the query.




Note: Use the Impala language reference to get information about syntax in addition to the SQL auto-complete feature that is built in. To view the language reference, click the book icon 📖 to the right of the query editor panel.

4. To run Hive queries:

- a) On the **Overview** page under Virtual Warehouses, click on the Hue button.

The Hive query editor is displayed:



- b) Click a database to view the tables it contains.

When you click a database, it sets it as the target of your query in the main query editor panel.

- c) Type a query in the editor panel and click the run icon ▶ to run the query.



Note: Use the Hive language reference to get information about syntax in addition to the SQL auto-complete feature that is built in. To view the language reference, click the book icon 📖 to the right of the query editor panel.

Viewing Hive query details

You can search Hive query history, compare two queries, download debug bundles for troubleshooting, and view query details, a graphical representation of the query execution plan, and DAG information on the Job Browser page in Hue.

Viewing Hive query history

The Queries tab on the Job Browser page in Hue displays all the queries that were run on all Hive Virtual Warehouses within a Database Catalog from various query interfaces, such as Beeline, Hive Warehouse Connector (HWC), Tableau, Hue, and other JDBC BI clients and tools.

About this task

Only Query Processor Administrators can view historical queries of all users to monitor resource utilization and control costs from the Hue Job Browser. Non-admin users can view only their queries.

Queries are retained in the backend database for 30 days by default, after which they are cleaned up. You can change the clean-up interval from the Database Catalog configurations.

Procedure

1. Go to the Cloudera Data Warehouse (CDW) web interface and open Hue from your Virtual Warehouse.
2. Click Jobs from the left assist panel.
The **Job Browser** page is displayed.
3. Click Queries.

The Hive queries that were run for the past seven days are displayed. You can select the time period for which you want to view the historical data.

You can also filter queries by their status.

Viewing Hive query information

The Query Info tab provides information such as, the Hive query ID, the user who executed the query, the start time, the end time, the total time taken to execute the query, the tables that were read and written, application ID, Directed Acyclic Graph (DAG) IDs, session ID, LLAP app ID, thread ID, and the queue against which the query was run.

Procedure

1. Go to the Cloudera Data Warehouse (CDW) web interface and open Hue from your Virtual Warehouse.
2. Click Jobs from the left assist panel.
The **Job Browser** page is displayed.

- Go to the **Queries** tab and click on the query for which you want to view the query details.

The following image shows the **Query Info** tab on the Hue web interface:

The screenshot displays the Hue web interface with the **Job Browser** header and **Jobs** / **Queries** tabs. The **Queries** tab is active, showing a table with query details. Below the table, the **Query Info** tab is selected, displaying the query text and various execution metrics.

QUERY ID	USER	STATUS
hive_20220509083516_a9d00c94-657a-4d80-9cc2-51851ec711eb		✓ SUCCESS

Buttons: **Queries**, **Kill**, **Download**

Query Info tab selected. Other tabs: Visual Explain, Timeline, Query Config, DAG Info, DAG Flow, DAG Swimlane, DAG Counters, DAG Configurations.

QUERY

```
SELECT
*
FROM
customer
WHERE
c_nationkey = 15
```

START TIME
3 minutes ago

END TIME
2 minutes ago

DURATION
50s

TABLES READ
customer (default)

TABLES WRITTEN
-

APPLICATION ID
application_1652085158072_0001

DAG ID
dag_1652085158072_0001_2

SESSION ID
76e59bed-40e6-4387-8c35-52606ecacaf4

LLAP APP ID

THREAD ID
HiveServer2-Background-Pool: Thread-297

QUEUE
None

Viewing explain plan for a Hive query

The Visual Explain feature provides a graphical representation of the query execution plan. The Explain plan is read from right to left. It provides details about every stage of query execution.

Procedure

- Go to the Cloudera Data Warehouse (CDW) web interface and open Hue from your Virtual Warehouse.
- Click **Jobs** from the left assist panel.
The **Job Browser** page is displayed.
- Go to the **Queries** tab and click on the query for which you want to view the query details.

4. Click on Visual Explain.

The following image shows the **Visual Explain** tab on the Hue web interface:



5. (Optional) Click to download the query explain plan in JSON format.

Viewing Hive query timeline

The Timeline tab provides a visual representation of Hive performance logs and shows the time taken by each stage of the query execution.

About this task

Following are the stages in which a query is executed:

- Pre-execution and DAG construction: It is the first phase of query execution and is executed on the Hive engine. It constitutes the time taken to compile, parse, and build the Directed Acyclic Graph (DAG) for the next phase of the query execution.
- DAG submission: It is the second phase in which the DAG that was generated in Hive is submitted to the Tez engine for execution.
- DAG runtime: It shows the time taken by the Tez engine to execute the DAG.
- Post-execution: It is the last phase of query execution in which the files in S3/ABFS are moved or renamed.

Duration data about each phase are distilled into more granular metrics based on query execution logs.

Procedure

1. Go to the Cloudera Data Warehouse (CDW) web interface and open Hue from your Virtual Warehouse.
2. Click Jobs from the left assist panel.
The **Job Browser** page is displayed.
3. Go to the **Queries** tab and click on the query for which you want to view the query details.
4. Click on Timeline.

The following image shows the **Timeline** tab on the Hue web interface:



Viewing configurations for a Hive query

The Query Config tab provides the configuration properties and settings that are used in a Hive query. You can use this tab to verify that configuration property values align with your expectations.

Procedure

1. Go to the Cloudera Data Warehouse (CDW) web interface and open Hue from your Virtual Warehouse.
2. Click Jobs from the left assist panel.
The **Job Browser** page is displayed.
3. Go to the **Queries** tab and click on the query for which you want to view the query details.
4. Click on Query Config.

The following image shows the **Query Config** tab on the Hue web interface:

Query Info	Visual Explain	Timeline	Query Config	DAG Info	DAG Flow	DAG Swimlane
Config Name		Config Value				
hadoop.security.group.mapping.ldap.posix.attr...		uidNumber				
dfs.block.invalidate.limit		1000				
yarn.admin.acl		*				
hive.repl.dump.metadata.only.for.external.table		true				
hive.exec.stagingdir		.hive-staging				
hive.druid.rollup		true				
yarn.federation.enabled		false				
yarn.app.mapreduce.am.job.committer.cancel...		60000				
hive.druid.broker.address.default		localhost:8082				
dfs.disk.balancer.max.disk.throughputInMBper...		10				
dfs.qjournal.select-input-streams.timeout.ms		20000				
hive.llap.io.orc.time.counters		true				
hive.repl.retain.prev.dump.dir		false				
hive.vectorized.execution.mapjoin.native.fast....		true				
dfs.provided.aliasmap.inmemory.leveldb.dir		/tmp				
yarn.nodemanager.process-kill-wait.ms		5000				
yarn.minicluster.use-rpc		false				
io.map.index.interval		128				

Viewing DAG information for a Hive query

Directed Acyclic Graph (DAG) is created by the Hive engine every time you query the Hive Virtual Warehouse. The Hive SQL queries are compiled and converted into a Tez execution graph also known as a DAG. DAG is a collection of vertices where each vertex executes a fragment of the query or script. Hue provides a web interface to view detailed information about DAGs.

About this task

Directed connections between vertices determine the order in which they are executed. For example, the vertex to read a table must be run before a filter can be applied to the rows of that table. As another example, consider a vertex that reads a user table that is very large and distributed across multiple computers and multiple racks. Reading the table is achieved by running many tasks in parallel.



Important: The DAG information tabs (**DAG Info**, **DAG Flow**, **DAG Swimlane**, **DAG Counters**, **DAG Configurations**) are displayed only if the Tez engine is used for query execution. The Tez engine is typically utilized for complex queries.

Procedure

1. Go to the Cloudera Data Warehouse (CDW) web interface and open Hue from your Virtual Warehouse.
2. Click Jobs from the left assist panel.

The **Job Browser** page is displayed.

3. Go to the **Queries** tab and click on the query for which you want to view the query details.
4. Click DAG Info to see the DAG ID, DAG name, the status of the query, the time taken to execute the DAG, start time, and end time.

The following image shows the **DAG Info** tab on the Hue web interface:

The screenshot shows the Hue web interface with the 'DAG Info' tab selected. At the top, there is a header bar with 'QUERY ID' (hive_20220509083516_a9d00c94-657a-4d80-9cc2-51851ec711eb), 'USER' (redacted), and 'STATUS' (SUCCESS). Below this is a navigation bar with tabs: Query Info, Visual Explain, Timeline, Query Config, DAG Info (selected), DAG Flow, DAG Swimlane, DAG Counters, and DAG Configurations. The main content area displays the following information:

DAG ID	dag_1652085158072_0001_2	DAG NAME	SELECT * FROM customer WHERE c_nationke...15 (Stage-1)
STATUS	SUCCEEDED	DURATION	00:00:50
START TIME	3 minutes ago	END TIME	2 minutes ago


The following table lists and describes the status of the Tez job:

Status	Description
Submitted	The DAG is submitted to Tez but is not running
Running	The DAG is currently running
Succeeded	The DAG was completed successfully
Failed	The DAG failed to complete successfully
Killed	The DAG was stopped manually
Error	An internal error occurred when executing the DAG

- Click DAG Flow to see the DAG in the form of a flowchart.

You can gain insight into the complexity and the progress of executing jobs, and investigate the vertices that have failures or are taking a long time to complete.

The following image shows the **DAG Flow** tab on the Hue web interface::




Here, the input to vertices Map 1 and Map 2 are the tables displayed in green boxes. Next, Map 2 depends on the result set generated by Map 1. Map 2 is the last vertex in the DAG flow and after it completes its execution, the query output is written to a file in a filesystem such as S3 or ABFS.

There are a few options to change the layout of the DAG flow. You can hide the input and the output nodes to view only the task vertices by clicking the Toggle source/sink visibility button. You can switch between the horizontal and vertical orientation by clicking the Toggle orientation button.

- Click DAG Swimlane to see the DAG of the vertices against time.

Each mapping and reducing task is a vertex. Each horizontal bar of the swimlane represents the total time taken by the vertex to complete the execution. The vertical lines indicate the time when the vertex was initialized, the time when the vertex started, the time when the first task started, the time when the last task was completed, and the time when the vertex finished its execution. When you mouse over the vertical line, the bubble displays the stage of the vertex execution and provides a timestamp. The vertical lines connecting two vertices denote the dependency of a vertex on another vertex.

The following image shows the **DAG Swimlane** tab on the Hue web interface:



In this example, Map 1 depends on the results of Map 5. Map 1 will finish its execution only when Map 5 finishes its execution successfully. Similarly, Reducer 2 depends on Map 1 to complete its execution.

The consolidated timeline shows the percentage of time each vertex took to complete executing.

- Click DAG Counters to see details such as the number of bytes read and written, number of tasks that initiated and ran successfully, amount of CPU and memory consumed, and so on.

The **DAG Counters** tab provides a way to measure the progress or the number of operations that occur within a generated DAG. Counters are used to gather statistics for quality control purposes or problem diagnosis.

The following image shows the **DAG Counters** tab on the Hue web interface:

Query Info	Visual Explain	Timeline	Query Config	DAG Info	DAG Flow	DAG Swimlane	DAG Counters	DAG Configurations
Group Name		Counter Name		DAG : dag_1652085158072_0001_2				
org.apache.tez.common.counters.DAGCounter		NUM_SUCCEEDED_TASKS		54				
org.apache.tez.common.counters.DAGCounter		TOTAL_LAUNCHED_TASKS		54				
org.apache.tez.common.counters.DAGCounter		DATA_LOCAL_TASKS		54				
org.apache.tez.common.counters.DAGCounter		AM_CPU_MILLISECONDS		5890				
org.apache.tez.common.counters.DAGCounter		AM_GC_TIME_MILLIS		22				
org.apache.tez.common.counters.FileSystem...		FILE_BYTES_WRITTEN		1074416910				
org.apache.tez.common.counters.FileSystem...		S3A_BYTES_READ		6742239795				
org.apache.tez.common.counters.FileSystem...		S3A_READ_OPS		1124				
org.apache.tez.common.counters.TaskCounter		TASK_DURATION_MILLIS		759357				
org.apache.tez.common.counters.TaskCounter		INPUT_RECORDS_PROCESSED		146519				
org.apache.tez.common.counters.TaskCounter		INPUT_SPLIT_LENGTH_BYTES		12387638515				
HIVE		CREATED_FILES		36				
HIVE		RECORDS_IN_Map_1		150000000				
HIVE		RECORDS_OUT_0		6003115				
HIVE		RECORDS_OUT_OPERATOR_FIL_5		6003115				
HIVE		RECORDS_OUT_OPERATOR_FS_7		6003115				
HIVE		RECORDS_OUT_OPERATOR_SEL_6		6003115				
HIVE		RECORDS_OUT_OPERATOR_TS_0		150000000				

8. Click DAG Configurations to see the Tez configuration details for a query that has a DAG associated with it. The following image shows the **DAG Configurations** tab on the Hue web interface:

Query Info		Visual Explain	Timeline	Query Config	DAG Info	DAG Flow	DAG Swimlane	DAG Counters	DAG Configurations
Config Name		DAG : dag_1612255715648_0000_1							
dfs.namenode.fs-limits.max-xattrs-per-inode		32							
dfs.namenode.delegation.token.always-use		false							
yam.nodemanager.runtime.linux.docker.delaye..		false							
yam.timeline-service.handler-thread-count		10							
yam.timeline-service.webapp.rest-csrf.custom..		X-XSRF-Header							
fs.s3a.retry.limit		7							
dfs.client.write.byte-array-manager.count-reset..		10000							
yam.nodemanager.linux-container-executor.cg..		/hadoop-yam							
mapreduce.shuffle.connection-keep-alive.time..		5							
mapreduce.client.libjars.wildcard		true							
hive.zookeeper.kerberos.enabled		false							

Terminating Hive queries

If a query is running for longer than expected, or you have accidentally triggered it, then you can stop the query to free up the resources. Hue also allows you to stop multiple queries at once.

About this task



Note: This feature is available only for Hive queries. Only admin users or Hue superusers can stop running queries.

Procedure

1. Go to the Cloudera Data Warehouse (CDW) web interface and open Hue from your Virtual Warehouse.
2. Click Jobs from the left assist panel.
The **Job Browser** page is displayed.
3. Go to the **Queries** tab.
A list of queries that were run is displayed.
4. Select the queries that you want to stop and click Kill.

Comparing Hive queries in Hue

You can compare two queries to know how each query is performing in terms of speed and cost-effectiveness. Hue compares various aspects of the two queries, based on which you can identify what changed between the executions of those two queries, and you can debug performance-related issues between different runs of the same query.

About this task

The query comparison report provides you a detailed side-by-side comparison of your queries.

For Hive queries, it includes recommendations for optimizing each query, metadata about the queries, visual explain for each query, query timeline, query configuration, Directed Acyclic Graph (DAG) information, DAG flows, DAG swimlanes, DAG counters, and DAG configurations.

Procedure

1. Go to the Cloudera Data Warehouse (CDW) web interface and open Hue from your Virtual Warehouse.
2. Click Jobs from the left assist panel.

The **Job Browser** page is displayed.

3. Go to the **Queries** tab.

A list of queries that were run is displayed.

4. Select the two queries you want to compare and click Compare.

Query comparison report for Hive queries:

QUERY ID	USER
hive_20220509083516_a9d00c94-657a-4d80-9cc2-51851ec711eb	...
hive_20220509083138_56c823bb-c635-4d1e-b5e4-b031b5c0e21e	...

Query Info	Visual Explain	Timeline	Query Config	DAG Info	DAG Flow	DAG Swimlane	DAG Counters	DAG Configurations
QUERY <pre>SELECT * FROM customer WHERE c_nationkey = 15</pre> START TIME 2 hours ago END TIME 2 hours ago DURATION 50s TABLES READ customer (default) TABLES WRITTEN - APPLICATION ID application_1652085158072_0001 DAG ID dag_1652085158072_0001_2 SESSION ID 76e59bed-40e6-4387-8c35-52606ecacaf4								

How to run a stored procedure from Hue in Cloudera Data Warehouse

HPL/SQL allows you to implement business logic using variables, expressions, flow-of-control statements, and iterations. HPL/SQL makes SQL-on-Hadoop more dynamic. You can leverage your existing procedural SQL skills, and use functions and statements to make your typical ETL development more productive. In Cloudera Data Warehouse, Hue provides a smart interface to run stored procedures.



Note: This feature is available only for Hive queries.

To run stored procedures from Hue, create a Hive Virtual Warehouse in CDW and enable the hplsql option in the hue-safety-valve field.

The following example creates a procedure and returns records by passing a cursor:

```
print 'Hello world';/
CREATE PROCEDURE greet(name STRING)
```

```

BEGIN
  PRINT 'Hello ' || name;
END;/
CREATE PROCEDURE even(cur OUT SYS_REFCURSOR)
BEGIN
  OPEN cur FOR
  SELECT n FROM NUMBERS
  WHERE MOD(n, 2) == 0;
END;/
CREATE PROCEDURE set_message(IN name STRING, OUT result STRING)
BEGIN
  SET result = 'Hello, ' || name || '!';
END;
-- Call the procedure and print the results
DECLARE str STRING;
CALL set_message('world', str);
PRINT str;

```



Attention: In the hplsql mode, you must terminate the commands using the forward slash (/). The semicolon (;) is used throughout procedure declarations and can no longer be relied upon to terminate a query in the editor.



Note: HPL/SQL does not support all types of Hive statements, such as JOIN or EXPLAIN. Refer to the [HPL/SQL Reference](#) for more information.

Enabling stored procedures for Hive in Cloudera Data Warehouse

To create, edit, and drop procedures and functions that are written in Hive Hybrid Procedural SQL (HPL/SQL) using the Hue query editor in CDW, you must enable the hplsql option in the hue-safety-valve field.

About this task



Important: Hue enables you to switch between Hive and HPL/SQL interpreters. By default, the regular Hive interpreter is enabled when you create a Hive Virtual Warehouse. To enable the HPL/SQL interpreter, you must update the configuration in the hue-safety-valve field in your Hive Virtual Warehouse. However, updating the hue-safety-valve overrides the default configuration. Therefore, to use both Hive and HPL/SQL interpreters, you must enable both by updating the configuration in the hue-safety-valve field.

Procedure

1. Log in to the Data Warehouse service as an administrator.
2. Go to Virtual Warehouse Edit CONFIGURATIONS Hue and select hue-safety-valve from the Configuration files drop-down list.
3. Add the following lines in the hue-safety-valve:

```

[notebook]
[[interpreters]]
[[hive]]
  name=Hive
  interface=hiveserver2
[[hplsql]]
  name=Hplsql
  interface=hiveserver2

```

4. Click APPLY.
5. Restart the Virtual Warehouse.

Enabling the SQL editor autocompleter

Autocompleter provides finely tuned SQL suggestions for Hive and Impala dialects while you enter queries into the editor window. See [Brand new Autocompleter for Hive and Impala](#) in the Hue blog.

About this task

Autocompleter is enabled by default. To manually enable or disable it, open the editor configuration panel and edit settings as follows:

Procedure

1. Log in to Hue and go to either the Hive or Impala editor.
2. Place your cursor in the editor window and then use one of the following keyboard shortcuts to open the editor configuration panel:

- On a Mac system, use the Command key followed by a hyphen and then a comma:

Command-,


- On a Windows system, use the Ctrl key followed by a hyphen and then a comma:

Ctrl-,



Tip: Type a question mark (?) anywhere but in the active editor window to open a menu of editor keyboard shortcuts.

3. To enable autocompletion, check the box adjacent to Enable Autocompleter. When you check Enable Autocompleter, Enable Live Autocompletion is automatically enabled as well. Place your cursor in the editor window to close the configuration panel.



4. To disable autocompletion:

- Uncheck Enable Live Autocompletion but leave Enable Autocompleter checked, and then place your cursor in the editor window to close the configuration panel. This disables live autocompletion, but if you want to use autocompletion while building your queries in the editor, enter the following key stroke sequence to activate autocompletion: Ctrl + Space Key
- Uncheck both Enable Autocompleter and Enable Live Autocompletion, and then click in the editor to close the configuration panel. This disables all autocompletion functionality.

Using governance-based data discovery

Hue can use the metadata tagging, indexing, and search features available in Apache Atlas data management. After integrating Hue with Atlas, classifications and indexed entities can be accessed and viewed in Hue. This topic shows you how to use metadata classifications in Hue.

Integration between Hue and Atlas is enabled by default, but if your administrator has disabled it, it must be re-enabled before you can use governance-based data discovery.

In Cloudera Data Warehouse, you can only view tags that are created in Atlas in Hue. You must create tags in Atlas.

Searching metadata tags

The SQL Editor in Hue provides a search text box where you can search on the metadata tags or classifications that are associated with your databases, tables, and columns.


About this task

You can search for tags or classifications in either the Hive or the Impala editors.

 **Note:** On clusters that use Apache Ranger for role-based access control, the Search mechanism does not display counts of popular values. Ranger ensures that Hue users can view only entities to which their user role (as configured and managed by Ranger) has been granted specific permissions.

Procedure

- 1. Go to Query Editor Impala or Hive.
- 2. To locate the tags or classifications in Apache Atlas, in the metadata search box located just to the right of the Query drop-down menu, type a tag: or classification: facet followed by its name. For example, type classification: wine as shown in the following image:



After you type the search facet and the tag or classification name in the search box, the `<DATABASE>.<TABLE>` where the tag or classification is found is returned. Click the `<DATABASE>.<TABLE>` to view the tags and classifications that have been defined for it.

Using Amazon S3 with Hue

Hue can read to and write to an Amazon S3 bucket.

 **Note:**

Using Azure Data Lake Storage Gen2 with Hue

Hue can read to and write to an Azure Data Lake Storage (ADLS) Gen2.



Note:

Only Hue superusers can view and access the ABFS file browser.

Granting permission to access S3 and ABFS File Browser in Hue

Only admin users can view and access S3 or ABFS File Browser in Hue after enabling it. You must manually grant application permissions to non-admin users and groups for them to be able to view and access S3 and ABFS File Browsers in Hue.

About this task

The following table lists the application permissions for each cloud storage type:

Cloud storage	Hue application permission
S3	filebrowser.s3_access:Access to S3 from filebrowser and filepicker.
ABFS (ADLS Gen2)	filebrowser.abfs_access:Access to ABFS from filebrowser and filepicker.
ADLS (Gen1)	filebrowser.adls_access:Access to ADLS from filebrowser and filepicker.
GS	Access to GS from filebrowser and filepicker.

Before you begin

You can only assign Hue application permissions to groups or users within a group. Add the users to a group to whom you want to grant access to the S3 or ABFS File Browsers.



Important: The "default" group in Hue does not have the permissions required to view S3 or ABFS File Browsers, by default. If your users belong to the default group, then you must manually grant the required permissions.

Procedure

1. Open Hue from Cloudera Data Warehouse Virtual Warehouse as an EnvironmentAdmin.
2. Go to admin Manage Users Groups .
3. Click on the group to whom you want to grant the filebrowser application permissions.
4. On the **Edit [***GROUP-NAME***]** page, select the required permission under the permission section and click Update group.

Supported non-ASCII and special characters in Hue

Auto-generated files may often introduce non-alphanumeric characters in the file and directory names that Hue does not support. This might cause the files or directories to not appear on the Hue File Browser. Review the list of non-alphanumeric, non-ASCII, diacritics (accents) characters supported in Hue for the following operations: upload, create, list in folder, view, and rename.

Table 1: Supported characters on HDFS

Special character symbol	Description	Filename support	Folder name support
~	Tilde	Fully supported	Fully supported

Special character symbol	Description	Filename support	Folder name support
@	Ampersat	Fully supported	Fully supported
#	Hash	Partially supported. Not supported for rename operations.	Unsupported
\$	Dollar sign	Fully supported	Fully supported
&	Ampersand	Fully supported	Fully supported
(	Left paranthesis	Fully supported	Fully supported
)	Right paranthesis	Fully supported	Fully supported
*	Asterisk	Fully supported	Fully supported
!	Exclamation mark	Fully supported	Fully supported
+	Plus	Fully supported	Fully supported
=	Equal	Fully supported	Fully supported
:	Colon	Unsupported	Unsupported
;	Semicolon	Fully supported	Fully supported
,	Comma	Fully supported	Fully supported
.	Period	Fully supported	Fully supported
?	Question mark Not supported with Knox.	Fully supported	Fully supported
/	Forward slash	Unsupported	Unsupported
\	Backslash	Unsupported	Unsupported
'	Apostrophe or single quote	Fully supported	Fully supported
T#####-ääö	Non-standard alphabets with diacritics and accents.	Fully supported	Fully supported

Table 2: Supported characters on S3

Special character symbol	Description	Filename support	Folder name support
~	Tilde	Fully supported	Fully supported
@	Ampersat	Fully supported	Fully supported
#	Hash	Partially supported. Not supported for view and rename operations.	Unsupported
\$	Dollar sign	Fully supported	Fully supported
&	Ampersand	Fully supported	Fully supported
(	Left paranthesis	Fully supported	Fully supported
)	Right paranthesis	Fully supported	Fully supported
*	Asterisk	Fully supported	Fully supported
!	Exclamation mark	Fully supported	Fully supported
+	Plus	Fully supported. Not supported for create operations on RAZ-enabled environments.	Fully supported
=	Equal	Fully supported	Fully supported
:	Colon	Unsupported	Unsupported
;	Semicolon	Fully supported	Fully supported
,	Comma	Fully supported	Fully supported

Special character symbol	Description	Filename support	Folder name support
.	Period	Fully supported	Fully supported
?	Question mark Not supported with Knox.	Fully supported	Partially supported. Not supported for list, upload, and view operations.
/	Forward slash	Unsupported	Unsupported
\	Backslash	Partially supported. Not supported for upload operations.	Partially supported. Not supported for upload operations.
'	Apostrophe or single quote	Fully supported	Fully supported
T#####-ääö	Non-standard alphabets with diacritics and accents.	Fully supported	Fully supported

Table 3: Supported characters on ABFS

Special character symbol	Description	Filename support	Folder name support
~	Tilde	Fully supported	Fully supported
@	Ampersat	Fully supported	Fully supported
#	Hash	Partially supported. Not supported for view and rename operations.	Unsupported
\$	Dollar sign	Fully supported	Fully supported
&	Ampersand	Fully supported	Fully supported
(	Left parenthesis	Fully supported	Fully supported
)	Right parenthesis	Fully supported	Fully supported
*	Asterisk	Fully supported	Fully supported
!	Exclamation mark	Fully supported	Fully supported
+	Plus	Fully supported	Fully supported
=	Equal	Fully supported	Fully supported
:	Colon	Unsupported	Unsupported
;	Semicolon	Fully supported	Fully supported
,	Comma	Fully supported	Fully supported
.	Period	Fully supported	Fully supported
?	Question mark Not supported with Knox.	Partially supported. Not supported for view and rename operations.	Partially supported. Not supported for list, rename, and view operations.
/	Forward slash	Unsupported	Unsupported
\	Backslash	Unsupported	Unsupported
'	Apostrophe or single quote	Fully supported	Fully supported
T#####-ääö	Non-standard alphabets with diacritics and accents.	Fully supported	Fully supported