

Cloudera Runtime 7.2.7

Using Hue

Date published: 2020-07-28

Date modified: 2021-02-04

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

Using Hue.....	4
Enabling the SQL editor autoCompleter.....	10
Using governance-based data discovery.....	11
Defining metadata tags.....	11
Searching metadata tags.....	15
Using Amazon S3 with Hue.....	16
Enabling S3 browser for Hue configured with IDBroker.....	16
Enabling S3 browser for Hue configured without IDBroker.....	17
Populating an S3 bucket.....	18
Creating a table from an Amazon S3 file.....	18
Exporting query results to Amazon S3.....	19
Troubleshooting errors when working with Amazon S3.....	22
Using Azure Data Lake Storage Gen2 with Hue.....	23
Enabling ABFS file browser for Hue configured with IDBroker.....	23
Enabling ABFS file browser for Hue configured without IDBroker.....	24
List of supported non-alphanumeric characters for file and directory names in Hue.....	25
Granting permission to access S3 and ABFS File Browser in Hue.....	26

Using Hue

Get started using Hue by analyzing and visualizing your data with Impala, a high-speed, low-latency SQL query engine.


About this task


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

Procedure


1. Download and unzip [one year of bike trips](#) from the Bay Area Bike Share program. This file is about 80 MB in size.

2. Create a table from the `~/babs_open_data_year_1/201402_babs_open_data/201402_trip_data.csv` file found in the unzipped `babs_open_data_year_1.zip` file:
 - a) In the Cloudera Manager Admin Console, select HueWebUIHue Load Balanced to launch Hue.
 - b) In the left navigation panel of Hue, make sure the default database is selected, and click the plus sign to create a table as shown in the following image:



If the default database is not selected, click the "less than" icon  that is next to the database icon in the left panel. This enables you to select the default database.

- c) In the center panel Importer UI, set Type to File.
- d) Drag the `201402_trip_data.csv` file to the Path field as shown in the following image:



- e) Set the formats as follows:
 - Field Separator = Comma (,)
 - Record Separator = New line
 - Quote Character = Double Quote

Then click Next at the bottom of the page.

- f) Set the properties Format = Text.
- g) Edit the FIELDS as follows:
 - Rename Bike # to Bike ID
 - Change the data type of ZipCode to string.
 - Remove all of the spaces in the Name fields.


Then click Submit at the bottom of the page.













3. Click Query at the top of the page and select EditorHive to open the Hive editor and then create a query.

- Enter the following query into the editor window:


```
SELECT * FROM default.201402_trip_data
LIMIT 10;
```

-

Click the execute icon  to run the query. The following rows are returned:

	Query History	Saved Queries	Results (10)																																																						
			<table><thead><tr><th></th><th>201402_trip_data.tripid</th><th>201402_trip_data.duration</th><th>201402_trip_data</th></tr></thead><tbody><tr><td></td><td>1</td><td>4576</td><td>63</td><td>8/29/2013 14:13</td></tr><tr><td></td><td>2</td><td>4607</td><td>70</td><td>8/29/2013 14:42</td></tr><tr><td></td><td>3</td><td>4130</td><td>71</td><td>8/29/2013 10:16</td></tr><tr><td></td><td>4</td><td>4251</td><td>77</td><td>8/29/2013 11:29</td></tr><tr><td></td><td>5</td><td>4299</td><td>83</td><td>8/29/2013 12:02</td></tr><tr><td></td><td>6</td><td>4927</td><td>103</td><td>8/29/2013 18:54</td></tr><tr><td></td><td>7</td><td>4500</td><td>109</td><td>8/29/2013 13:25</td></tr><tr><td></td><td>8</td><td>4563</td><td>111</td><td>8/29/2013 14:02</td></tr><tr><td></td><td>9</td><td>4760</td><td>113</td><td>8/29/2013 17:01</td></tr><tr><td></td><td>10</td><td>4258</td><td>114</td><td>8/29/2013 11:33</td></tr></tbody></table>		201402_trip_data.tripid	201402_trip_data.duration	201402_trip_data		1	4576	63	8/29/2013 14:13		2	4607	70	8/29/2013 14:42		3	4130	71	8/29/2013 10:16		4	4251	77	8/29/2013 11:29		5	4299	83	8/29/2013 12:02		6	4927	103	8/29/2013 18:54		7	4500	109	8/29/2013 13:25		8	4563	111	8/29/2013 14:02		9	4760	113	8/29/2013 17:01		10	4258	114	8/29/2013 11:33
	201402_trip_data.tripid	201402_trip_data.duration	201402_trip_data																																																						
	1	4576	63	8/29/2013 14:13																																																					
	2	4607	70	8/29/2013 14:42																																																					
	3	4130	71	8/29/2013 10:16																																																					
	4	4251	77	8/29/2013 11:29																																																					
	5	4299	83	8/29/2013 12:02																																																					
	6	4927	103	8/29/2013 18:54																																																					
	7	4500	109	8/29/2013 13:25																																																					
	8	4563	111	8/29/2013 14:02																																																					
	9	4760	113	8/29/2013 17:01																																																					
	10	4258	114	8/29/2013 11:33																																																					

4. Click Query at the top of the page and select EditorImpala to open the Impala SQL editor and then create a query.
 - a. In the left panel, click the refresh icon and select Perform incremental metadata update to make the new table visible to Impala:




The screenshot shows the Hue interface with a dialog box open over the 'Tables' panel. The dialog box has three radio button options: 'Clear cache', 'Perform incremental metadata update.', and 'Invalidate all metadata and rebuild index.'. The 'Perform incremental metadata update.' option is selected and highlighted with a yellow box. Below it, the text reads 'This will sync missing tables.'. The 'Invalidate all metadata and rebuild index.' option has a warning below it: 'WARNING: This can be both resource and time-intensive.'. A blue 'Refresh' button is at the bottom right of the dialog. In the background, the 'Tables' panel shows a table named '201402_trip_data' with columns: tripid (bigint), duration (bigint), startdate (string), startstation (string), startterminal (bigint), enddate (string), endstation (string), endterminal (bigint), bikeid (string), subscriptiontype (string), and zipcode (string). A yellow arrow points to the refresh icon (a circular arrow) next to the table name, with the text 'refresh icon' written in yellow below it.

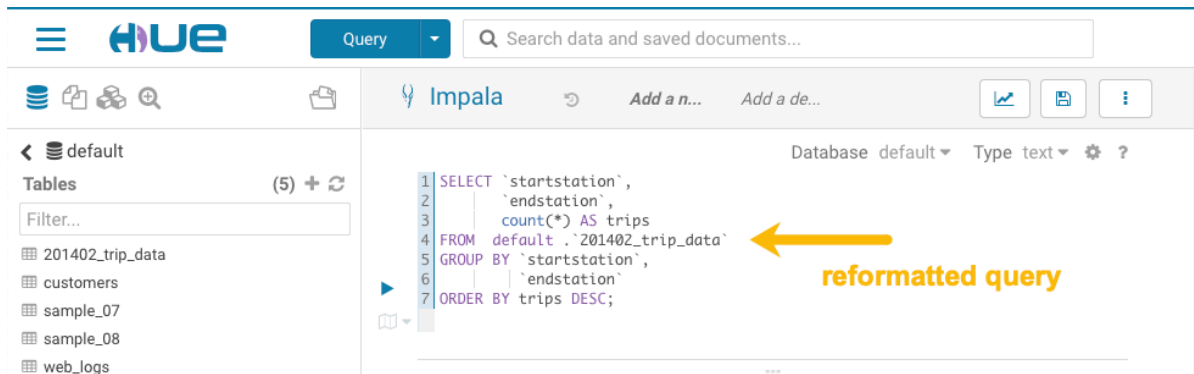
- b. Enter the following query into the editor window:

```
select 'startstation', 'endstation', count(*) as trips from default.'201402_trip_data'
group by 'startstation', 'endstation' order by trips desc;
```


c. Click the down arrow just under the execution icon and select Format:



This reformats the query:



d.

Click the save icon , enter a query name, and click Save.

e.

Click the execute icon  to run the query.

5. Create a bar chart that is based on the query results:

a.



Click the chart icon and then select Bars.

Impala interface showing a query and its results. The query is:

```

1 SELECT `startstation`,
2        `endstation`,
3        count(*) AS trips
4 FROM default.`201402_trip_data`
5 GROUP BY `startstation`,
6          `endstation`
7 ORDER BY trips DESC;
    
```

The results table shows the top 10 stations by trip count:

startstation	endstation
1	Harry Bridges Plaza (Ferry Building) - Embarcadero at Sansome
2	Townsend at 7th - San Francisco Caltrain (Townsend at 4th)
3	San Francisco Caltrain 2 (330 Townsend) - Townsend at 7th
4	Market at Sansome - 2nd at South Park
5	Embarcadero at Sansome - Steuart at Market
6	2nd at South Park - Market at Sansome
7	San Francisco Caltrain (Townsend at 4th) - Harry Bridges Plaza (Ferry Building)
8	2nd at Townsend - Harry Bridges Plaza (Ferry Building)
9	2nd at Townsend - Harry Bridges Plaza (Ferry Building)
10	2nd at Townsend - Harry Bridges Plaza (Ferry Building)


The 'Columns (4)' panel shows 'startstation' selected, and the 'Charts' menu has 'Bars' highlighted with a yellow arrow.

b. Set the bar chart elements as follows:

- X-AXIS = startstation
- Y-AXIS = trips
- LIMIT = 10

Configuration panel for the bar chart:

- TYPE: Bars
- X-AXIS: startstation
- Y-AXIS: trips
- GROUP: Choose a column to pivo...
- LIMIT: 10
- SORTING: [Icons for sorting]




6.

Create a pie chart by clicking the chart icon again



and then select Pie.

7.

Download the query results by clicking the download icon  and selecting in what format you want to download, copy, or export the results.

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-,

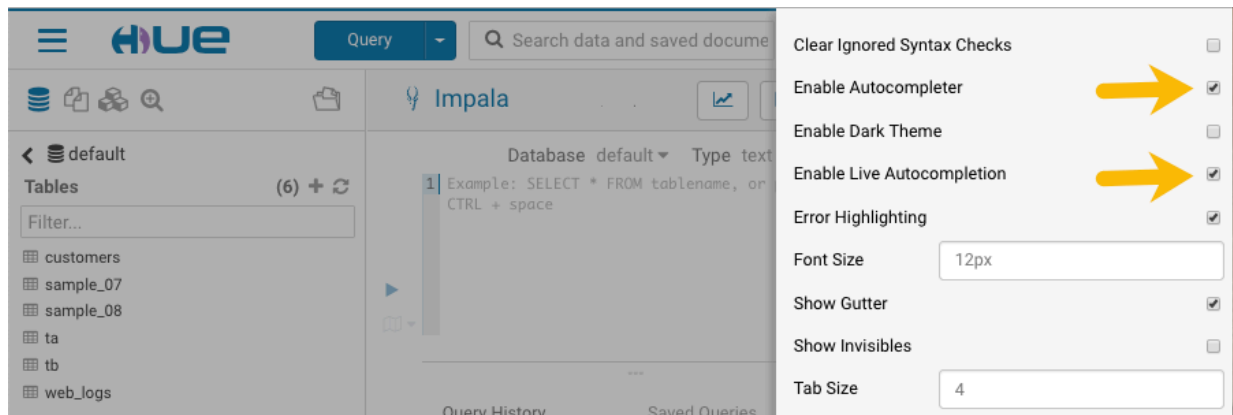
- On a Windows system, use the Ctrl key followed by a hyphen and then a comma:

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.

Defining metadata tags

The first step to access and use metadata that is stored in Cloudera Navigator is to define metadata tags in Hue. Then you can use the search and indexing capabilities from Cloudera Navigator in Hue and add metadata to Cloudera Navigator without leaving the Hue UI.

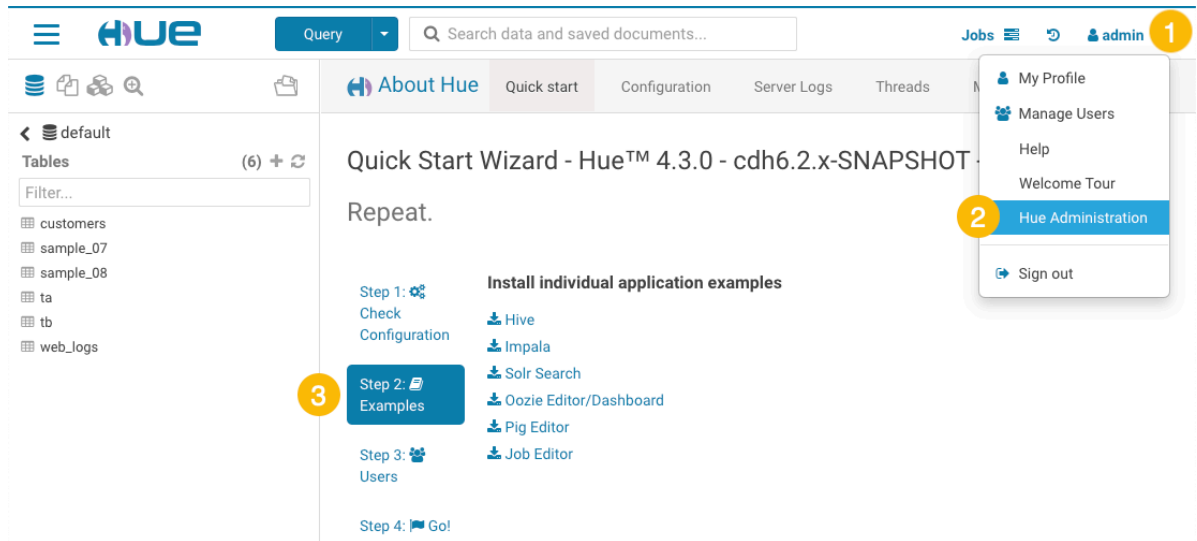
About this task

As the Hue administrator, install sample tables, and then refresh the metadata. This example uses Impala.


Procedure

1. Prepare Tables for Tagging

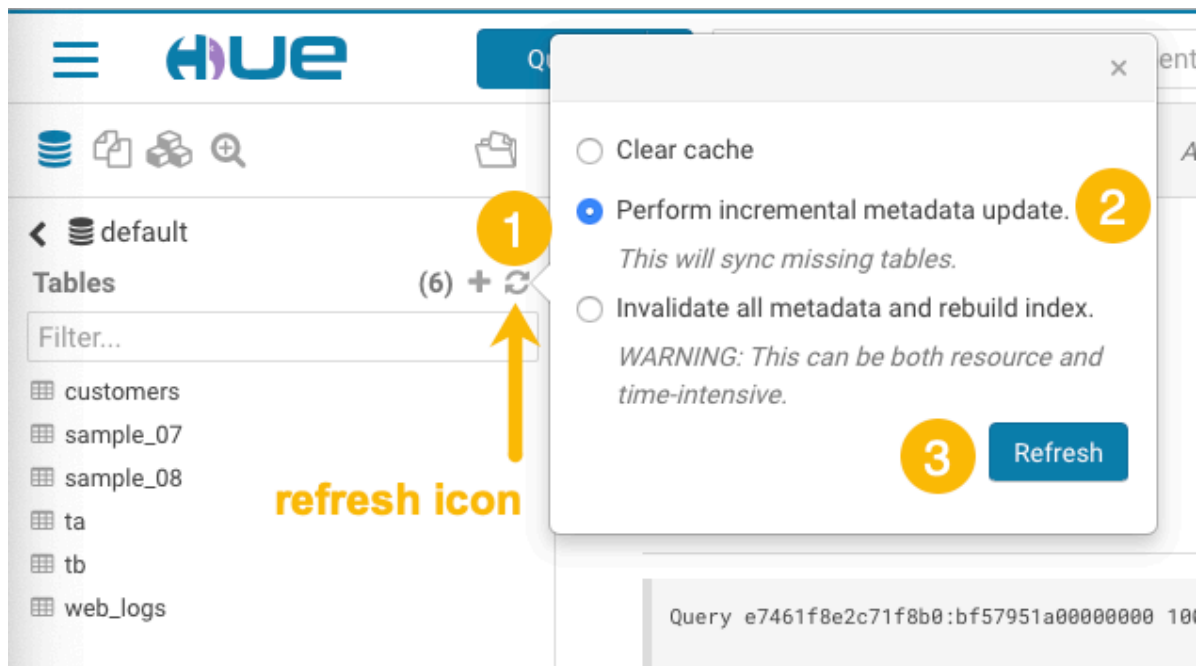
- a. In the Cloudera Manager Admin Console, launch Hue by selecting ClustersHueHue Web UIHue Load Balanced.
- b. In the top right corner of the page, select <USER_NAME>Hue AdministrationStep 2: Examples



c.

Install the sample tables for Hive and Impala by clicking the download icons .

- d. At the top of the page, click Query and then select EditorImpala to go to the Impala SQL editor, click the refresh icon, select Perform incremental metadata update, and then click Refresh:




e.

In the upper left corner of the page, select the menu icon  and then select BrowsersTables to load the Table Browser.

- f.** In the Table Browser, click Refresh on the right side of the page. This ensures that the sample tables are available to work with in Hue.

2. Tag a Database, Table, and Field

a.

At the top of the Table Browser page, click the table icon  to display the Databases page, and then click default or some other database. This loads a detail view of the database

b. To add a database tag, click Add tags..., enter a tag name in the text box, click Add <TAG_NAME>..., and then click the check on the right to save the tag:




1 Enter a tag name.

2 Click "Add <tag_name>..."

3 Click the check to save the tag.

c. To add a table tag, click a table name, for example customers, which loads a detail view of the table where you can define a tag:

1. Click Add tags...:



Click

2. Enter a tag name in the text box, click Add <TAG_NAME>..., and then click the check on the right to save the tag. This is the same process you performed for tagging the database in Step 2.

d. To add a field tag:

1. Click the information icon by a column name:

Impala ▾

Databases > default > customers

Add a description...

test_table_tag

Add properties...

Overview Sample (53) Details Privileges

PROPERTIES **STATS** ↻

Table Files 1 Rows 0 Total size 15.44 KB

Managed and stored in location Data last updated on 7:02 AM

Created by admin on 7:02 AM

SCHEMA

Filter...

Column (5)	Type	Description
<i>i</i> id	int	Add a description...
<i>i</i> name	string	Add a description...
<i>i</i> email_preferences	struct	Add a description...

2. In the dialog box that opens, click Add tags..., enter a tag name in the text box, click Add <TAG_NAME>..., and then click the check on the right to save the tag.

Now, you have added tags to a database, a table, and a field (column) and you can search using these tag names.

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.

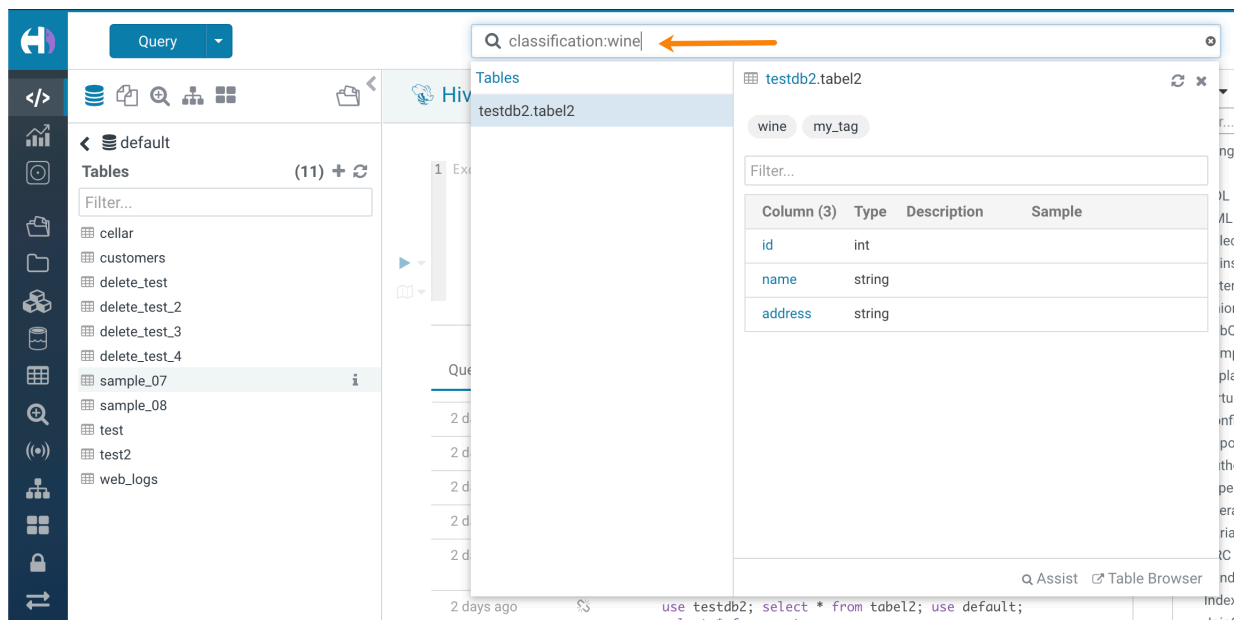


Note: On clusters that use Apache Ranger for role-based access control, the Search mechanism does not display counts of popular values. Ranger ensures that Hue users can view only entities to which their user role (as configured and managed by Ranger) has been granted specific permissions.

Procedure

1. Go to Query Editor Impala or Hive.

- 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: You can use the S3 browser to store and retrieve data from Amazon S3 buckets into Hue and vice versa. If your CDP cluster is secured using Knox, then the Hue users will be able to access the S3 folders as per the access rules set on the individual folders. The access permissions are governed by IDBroker.

Only Hue superusers can view and access the S3 browser.

Enabling S3 browser for Hue configured with IDBroker

You can access the S3 buckets from Hue to upload files and tables to S3 and import CSV files as tables directly in Hue by enabling the S3 browser on the Hue user interface.

Before you begin

If you have set up authentication using Knox IDBroker on your cluster, then Hue automatically detects and uses the IDBroker mappings from your cluster's `core-site.xml` file. Verify that the following property is present in the `core-site.xml` file:

```
<property>
  <name>fs.s3a.ext.cab.address</name>
  <value>https://<IDBROKERURL>:8444/gateway</value>
</property>
```

This property is automatically appended to the `core-site.xml` file when you enable Knox IDBroker on your cluster.

Procedure

1. Sign in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration .
3. Enter the following in the Hue Server Advanced Configuration Snippet (Safety Valve) for hue_safety_valve_server.ini field:

```
[desktop]
# Remove the file browser from the blacklisted apps.
# Tweak the app_blacklist property to suit your app configuration.
app_blacklist=spark,zookeeper,hive,hbase,search,oozie,jobsub,pig,sqoop,security
[aws]
has_iam_detection=true
[[aws_accounts]]
[[[default]]]
region=[**AWS-REGION**]
# Set a particular S3 bucket as the default
[filebrowser]
remote_storage_home=s3a://[**S3-BUCKET-NAME**]
```

The custom configuration is stored in the hue_safety_valve_server.ini file.

4. Click Save Changes.
5. Restart the Hue service.

The S3 file browser icon appears on the left Assist panel as well as on the left navigation bar on the Hue web interface.

What to do next

You must manually grant the following application permission to non-admin users and groups for them to be able to view and access S3 File Browser in Hue: filebrowser.s3_access:Access to S3 from filebrowser and filepicker..

You must also add the CDP users and groups to IAM role mappings.

Related Information

[Granting permission to access S3 and ABFS File Browser in Hue](#)

Enabling S3 browser for Hue configured without IDBroker

You can access the S3 buckets from Hue to upload files and tables to S3 and import CSV files as tables directly in Hue by enabling the S3 browser. You must specify the AWS access key and secret access key information along with the AWS region in the Hue Server Advanced Configuration Snippet if you have not set up authentication using Knox IDBroker.

Procedure

1. Sign in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration .
3. Enter the following in the Hue Server Advanced Configuration Snippet (Safety Valve) for hue_safety_valve_server.ini field:

```
[desktop]
# Remove the file browser from the blacklisted apps.
# Tweak the app_blacklist property to suit your app configuration.
app_blacklist=spark,zookeeper,hive,hbase,search,oozie,jobsub,pig,sqoop,security
[aws]
[[aws_accounts]]
```

```
[[[default]]]
access_key_id=[**AWS-ACCESS-KEY**]
secret_access_key=[**SECRET-ACCESS-KEY**]
region=[**AWS-REGION**]

# Set a particular S3 bucket as the default
[filebrowser]
remote_storage_home=s3a://[**S3-BUCKET-NAME**]
```

The custom configuration is stored in the hue_safety_valve_server.ini file.

4. Click Save Changes.
5. Restart the Hue service.

The S3 file browser icon appears on the left Assist panel as well as on the left navigation bar on the Hue web interface.

Related Information

[Granting permission to access S3 and ABFS File Browser in Hue](#)

Populating an S3 bucket

Use the Hue Web UI to populate buckets in Amazon S3.

About this task

Use open data from the U.S. Geological Survey to demonstrate how to populate and S3 bucket with Hue.

Procedure

1. Download [30 days of earthquake data](#) (all_month.csv) from the [USGS](#) (~2 MB).
2. Log on to the Hue Web UI from Cloudera Manager.
3. Select File BrowserS3 Browser.
4. Click NewBucket, name it "quakes_<any unique id>" and click Create.



Tip: Unique bucket names are important per S3 [bucket naming conventions](#).

5. Navigate into the bucket by clicking the bucket name
6. Click NewDirectory, name it "input" and click Create.
7. Navigate into the directory by clicking the directory name.
8. Click Upload and select, or drag, all_month.csv. The path is s3a://quakes/input/all_month.csv.




Important: Do not add anything else to the "input" directory—no extra files, no directories.

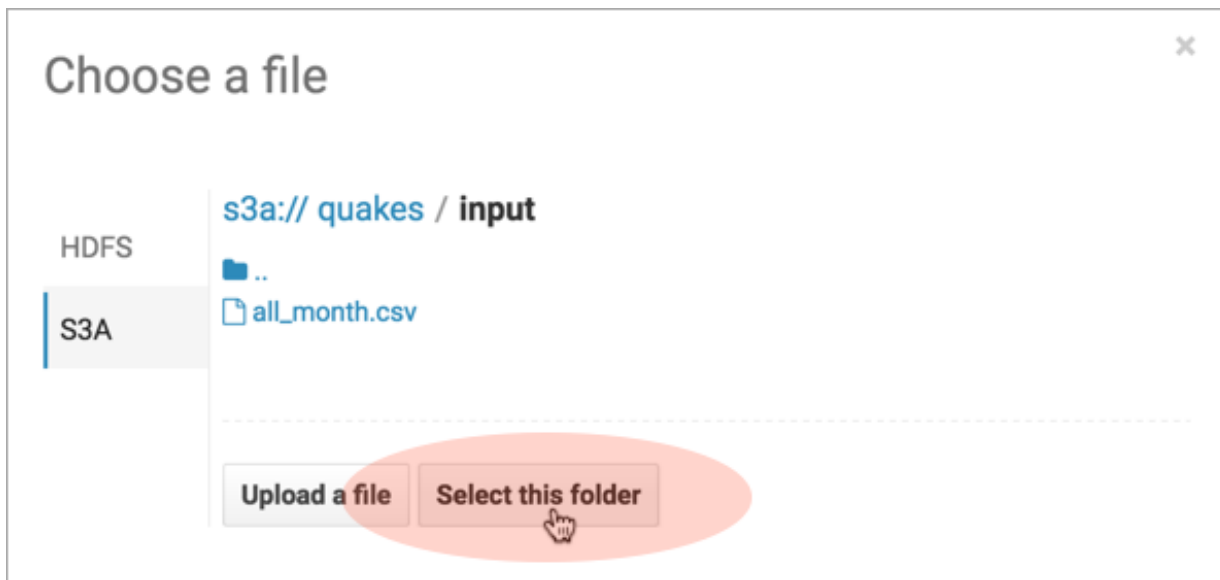
Creating a table from an Amazon S3 file

Using Hue to create a table from an Amazon S3 file streamlines the process.

Procedure

1. Go to the Metastore Manager by clicking Data BrowsersMetastore Tables.
2. Create a new table from a file by clicking .
3. Enter a Table Name such as "earthquakes".

4. Browse for the Input Directory, `s3a://quakes/input/`, and click Select this folder.



5. Select Create External Table from the Load Data menu and click Next.
6. Delimit by Comma(,) and click Next.
7. Click Create Table.
8. Click the Browse Data icon to automatically generate a SELECT query in the Hive editor:

```
SELECT * FROM `default`.`earthquakes` LIMIT 10000;
```


Exporting query results to Amazon S3

Use Hue to export query results to Amazon S3 as a custom file, a MapReduce file, or as a table.

Procedure

1. Run and Export Results in Hive

a)

Run the query by clicking Execute 

b)

Click Get Results 

c) Select Export to open the Save query result dialog.



2. Save Results as Custom File

a) Select In store (max 10000000 cells) and open the Path to CSV file dialog.


b) Navigate into the bucket, s3a://quakes.

c) Create folder named, "output."

d) Navigate into the output directory and click Select this folder.


e) Append a file name to the path, such as quakes.csv.

f) Click Save. The results are saved as s3a://quakes/output/quakes.csv.



3. Save Results as MapReduce files

- a) Select In store (large result) and open the Path to empty directory dialog.
- b) Navigate into the bucket, s3a://quakes.
- c) If you have not done so, create a folder named, "output."
- d) Navigate into the output directory and click Select this folder.
- e) Click Save. A MapReduce job is run and results are stored in s3a://quakes/output/.


**4. Save Results as Table**

- a) Run a query for "moment" earthquakes and export:

```
SELECT time, latitude, longitude, mag
FROM `default`.`earthquakes`
WHERE magtype IN ( 'mw', 'mwb', 'mwc', 'mwr', 'mww' );
```




- b) Select A new table and input <database>.<new table name>.
- c) Click Save.
- d)

Click Browse Data  to view the new table.



Troubleshooting errors when working with Amazon S3

When you use Hue with Amazon S3, you might encounter these errors. To view buckets, directories, and files that have been added to your S3 account, restart the Hue service.

Error Message	Example	Possible Solution
Failed to access path:	<pre>Failed to access path: "s3a:// quakes". Check that you have access to read this bucket and that the region is correct.</pre>	<p>Check your bucket region:</p> <ol style="list-style-type: none"> 1. Log on to your AWS account and navigate to the S3 service. 2. Select your bucket, for example "quakes", and click Properties. 3. Find your region. If it says US Standard, then region=us-east-1. 4. Update your configuration in Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini. 5. Save your changes and restart Hue.
The table could not be created.	<pre>The table could not be created. Error while compiling statement: FAILED: SemanticException: com.cloudera.com.amazonaws.AmazonClientException: Unable to load AWS credentials from any provider in the chain.</pre>	<p>Set your S3 credentials in Hive core-site.xml:</p> <ol style="list-style-type: none"> 1. In Cloudera Manager, go to HiveConfiguration. 2. Filter by CategoryAdvanced. 3. Set your credentials in Hive Service Advanced Configuration Snippet (Safety Valve) for core-site.xml. <ol style="list-style-type: none"> a. Click the  button and input Name and Value for fs.s3a.AccessKeyId. b. Click the  button and input Name and Value for fs.s3a.SecretAccessKey. 4. Save your changes and restart Hive.
The target path is a directory.	-	Remove any directories or files that may have been added to s3a://quakes/input/ (so that all_month.csv is alone).
Bad status for request TFetchResultsReq ... Not a file	<pre>Bad status for request TFetchResultsReq(TFetchResultsReq((status=TStatus orCode=0, errorMessage='java.io.IOException: java.io.IOException: Not a file: s3a://Not a file: s3a:// quakes/ input/ output' ...</pre>	<p>Remove any directories or files that may have been added to s3a://quakes/input/ (so that all_month.csv is alone). Here, Hive cannot successfully query the earthquakes table (based on all_month.csv) due to the directory, s3a://quakes/input/output.</p> <p> Tip: Run tail -f against the Hive server log in: /var/log/hive/.</p>

Using Azure Data Lake Storage Gen2 with Hue

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



Note: You can use the ABFS file browser to store and retrieve data from ADLS Gen2 into Hue and vice versa. If your CDP cluster is secured using Knox, then the Hue users will be able to access the ADLS as per the access rules set on the individual components. The access permissions are governed by IDBroker.

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

Enabling ABFS file browser for Hue configured with IDBroker

You can directly access the Azure file storage from Hue to upload files and tables to Azure file storage and import CSV files as tables directly in Hue by enabling the ABFS file browser on the Hue user interface.

Before you begin

If you have set up authentication using Knox IDBroker on your cluster, then Hue automatically detects and uses the IDBroker mappings from your cluster's core-site.xml file. Verify that the following property is present in the core-site.xml file:

```
<property>
  <name>fs.azure.ext.cab.address</name>
  <value>https://<IDBROKERURL>:8444/gateway</value>
</property>
```

This property is automatically appended to the core-site.xml file when you enable Knox IDBroker on your cluster.

You can also check whether there is a mapped role for the group(s) associated with the authenticated user by going to Management Console Environments Actions Manage Access IDBroker Mappings . If no role is mapped, then you may see the following error: Failed to obtain storage credentials from IDBroker with error: 403 Client Error: Forbidden for url: <url> { "error": "There is no mapped role for the group(s) associated with the authenticated user.", "auth_id": "<user-id>" }. To resolve this, add the user and the Data Access role to the current mappings on the **IDBroker Mappings** page.

Procedure

1. Sign in to Cloudera Manager as an Admin user.
2. Go to Clusters Hue service Configuration .
3. Specify the following in the Hue Server Advanced Configuration Snippet (Safety Valve) for hue_safety_valve_server.ini field:

```
[desktop]
# Remove the file browser from the blacklisted apps.
# Tweak the app_blacklist property to suit your app configuration.
app_blacklist=spark,zookeeper,hive,hbase,search,oozie,jobsub,pig,sqoop,security
[azure]
[[abfs_clusters]]
[[[default]]]
  fs_defaultfs=abfs://<CONTAINER_NAME>@<STORAGE_NAME>.dfs.core.windows.net
  webhdfs_url=https://<STORAGE_NAME>.dfs.core.windows.net/
```

The custom configuration is stored in the hue_safety_valve_server.ini file.

- (Optional) If the Knox server and Hue are not present on the same host, then add the Hue server's FQDN to the trusted origins property in the Hue Server Advanced Configuration Snippet (Safety Valve) for hue_safety_valve_server.ini field:

```
[desktop]
[[session]]
trusted_origins=[***HUE-SERVER-FQDN***]
```

- Click Save Changes.

- Restart the Hue service.

The ABFS file browser icon appears on the left Assist panel as well as on the left navigation bar on the Hue web interface.

What to do next

You must manually grant the following application permission to non-admin users and groups for them to be able to view and access ABFS File Browser in Hue: filebrowser.abfs_access:Access to ABFS from filebrowser and file picker.

You must also add the CDP users and groups to IAM role mappings.

Related Information

[Granting permission to access S3 and ABFS File Browser in Hue](#)

Enabling ABFS file browser for Hue configured without IDBroker

You can directly access the Azure file storage from Hue to upload files and tables to Azure file storage and import CSV files as tables directly in Hue by enabling the ABFS file browser on the Hue user interface. You must specify the Azure client ID and the client secret ID information along with the tenant ID in the Hue Server Advanced Configuration Snippet if you have not set up authentication using Knox IDBroker.

Procedure

- Sign in to Cloudera Manager as an Admin user.
- Go to Clusters Hue service Configuration .
- Specify the following in the Hue Server Advanced Configuration Snippet (Safety Valve) for hue_safety_valve_server.ini field:

```
[desktop]
# Remove the file browser from the blacklisted apps.
# Tweak the app_blacklist property to suit your app configuration.
app_blacklist=spark,zookeeper,hive,hbase,search,oozie,jobsub,pig,sqoop,security
[azure]
  [[azure_accounts]]
    [[[default]]]
      client_id=<CLIENT_ID>
      client_secret=<CLIENT_SECRET_ID>
      tenant_id=<TENANT_ID>

  [[abfs_clusters]]
    [[[default]]]
      fs_defaultfs=abfs://<CONTAINER_NAME>@<STORAGE_NAME>.dfs.core.windows.net
      webhdfs_url=https://<STORAGE_NAME>.dfs.core.windows.net/
```

The custom configuration is stored in the hue_safety_valve_server.ini file.

- Click Save Changes.

5. Restart the Hue service.

The ABFS file browser icon appears on the left Assist panel as well as on the left navigation bar on the Hue web interface.

Related Information

[Granting permission to access S3 and ABFS File Browser in Hue](#)

List of supported non-alphanumeric characters for file and directory names in Hue

Auto-generated files may often introduce non-alphanumeric characters in the filenames which are not supported by Hue. This might cause the files or directories to not appear on the Hue File Browser. Review the list of non-alphanumeric characters supported in Hue to avoid running into this issue.

The following table lists the supported non-alphanumeric characters in Hue:

Table 1: Non-alphanumeric characters supported in Hue

Special character symbol	Description
~	Tilde
@	Ampersat
#	Hash
\$	Dollar sign
&	Ampersand
(Left paranthesis
)	Right paranthesis
*	Asterisk
!	Exclamation mark
+	Plus
=	Equal
:	Colon Not supported with Knox.
;	Semicolon
,	Comma
.	Period
?	Question mark Not supported with Knox.
/	Forward slash Not supported with Knox.
\	Backslash
'	Apostrophe or single quote

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. Sign in to Hue as an Administrator.
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.