

Cloudera Runtime 7.2.18

Administering Hue

Date published: 2020-08-10

Date modified: 2024-02-29

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

Reference architecture.....	5
Hue configuration files.....	6
Hue configurations in CDP Runtime.....	7
Hue Advanced Configuration Snippet.....	9
Advanced safety valves.....	9
Environment safety valves.....	9
Set up SQL AI Assistant.....	10
Preparing the Microsoft Azure OpenAI service.....	10
Preparing the Amazon Bedrock Service.....	11
Preparing the OpenAI platform.....	11
Configuring the Hue SQL AI Assistant.....	12
Service and model-related configurations for setting up the Hue SQL AI Assistant.....	13
Hue logs.....	14
Standard stream logs.....	14
Hue service Django logs.....	15
Enabling DEBUG logging for Hue logs.....	16
Enabling httpd log rotation for Hue.....	16
Hue supported browsers.....	17
Enabling cache-control HTTP headers when using Hue.....	17
Setting up a Hue service account with a custom name.....	18
To validate newly configured custom user and group.....	19
Options to restart the Hue service.....	20
Customizing the Hue web interface.....	21
Adding a custom banner in Hue.....	21
Changing the page logo in Hue.....	21
Adding a splash screen in Hue.....	24
Setting the cache timeout.....	24
Enabling or disabling anonymous usage date collection.....	25

Disabling the share option in Hue.....	25
Enabling Hue applications with Cloudera Manager.....	26
Running shell commands.....	26
Downloading and exporting data from Hue.....	28
Enabling a multi-threaded environment for Hue.....	29
Adding Query Processor service to a cluster.....	29
Removing Query Processor service from cluster.....	31
Enabling the Query Processor service in Hue.....	31
Adding Query Processor admin users and groups.....	32
Cleaning up old queries.....	32
Downloading debug bundles.....	33
Configuring Hue to handle HS2 failover.....	34
Enabling Spark 3 engine in Hue.....	35
Using Hue scripts.....	36
Configurations for submitting a Hive query to a dedicated queue.....	37
Configuring timezone for Hue.....	37

Reference architecture

Hue server can support approximately 25 concurrent users, depending on what tasks the users are performing. Most scaling issues occur as a result of users performing resource-intensive operations and not from the number of users. For example, large downloads of query results can impact resource availability for the other users who are using the same Hue instance during the download operation. During that time, the users can experience slow performance. Another common cause of noticeable performance changes are slow RPC calls between Hue and another service. When this happens, queries may appear to suddenly "stop responding" after they are submitted.

As a guide, 2 Hue servers can support up to:

- 100 unique users per week
- 50 users per hour at peak times executing up to 100 queries

A typical setup is 2 Hue servers.

General guidelines

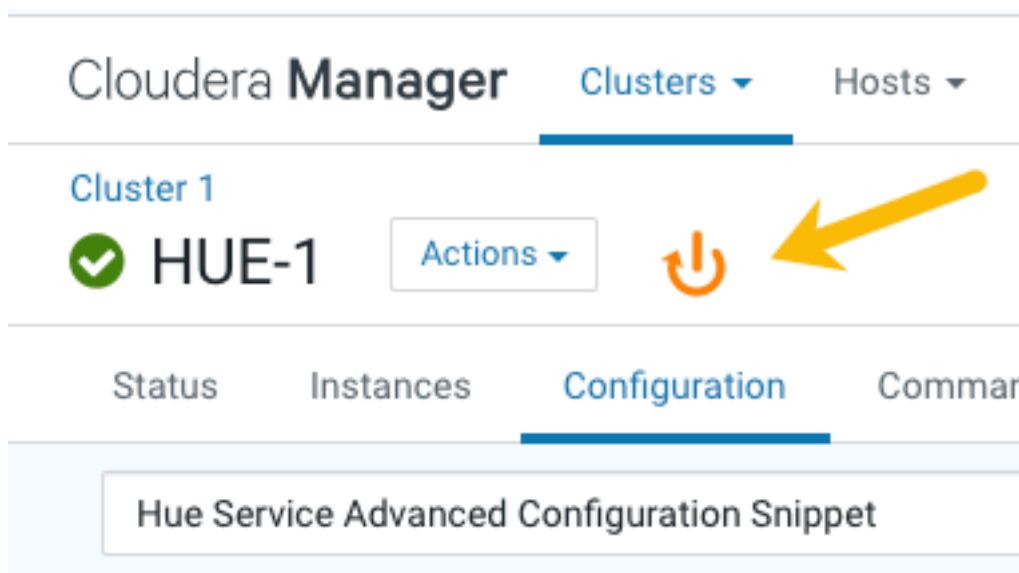
- Deploy a load balancer in front of Hue.
- Use a production-quality database.
- Ensure that other services, such as Impala, Hive, and Oozie, are healthy and not impacted by too few resources. If these services are hanging, it adversely affects Hue performance.
- Consider moving workloads that are subject to SLAs (service-level agreements) or considered "noisy neighbors" to their own compute cluster. Noisy neighbors are workloads that use the majority of available resources and cause performance issues.
- Limit the number of rows that are returned for queries.

One way to limit the number of rows returned is to specify a value for the `download_row_limit` configuration property for the Hue Beeswax application. This property can be set in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` property in Cloudera Manager:

1. In Cloudera Manager, click `HueConfiguration`, and enter Hue Service Advanced Configuration Snippet in the search text box.
2. In the text box for the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini`, add the following configuration information:

```
[beeswax]
download_row_limit=number_of_rows
```

3. Click Save Changes and click the restart icon at the top of the page to restart the Hue service:



Hue configuration files

Hue roles are configured with the following three configuration files: `hue.ini`, `hue_safety_valve.ini`, and `hue_safety_valve_server.ini`.

The `hue.ini` file is the first file that is auto-generated when you add the Hue service to your CDP cluster using Cloudera Manager. The `hue_safety_valve.ini` file is used to override bugs or configure properties that are missing from the Cloudera Manager UI. The `hue_safety_valve_server.ini` file is used to configure only the Hue role, and not the full Hue service (which includes the Hue Load Balancer). The `hue_safety_valve_server.ini` file is not used in practice.

Cloudera recommends that you do not edit the `.ini` files from the command line because they are stored in dynamic directories named by process ID and populated from the Cloudera Manager database. To add configurations that you cannot add directly from the Cloudera Manager UI, such as Authentication Backend for SAML, use the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field under the Hue service configurations in Cloudera Manager.

Run the following command to view the `.ini` files per process ID:

```
ls -ltr /var/run/cloudera-scm-agent/process/`ls -valrt /var/run/cloudera-scm-agent/process | grep HUE_SERVER | tail -1 | awk '{print $9}'`
```

Figure 1: Terminal showing Hue configuration files

```
[root@hue4-cdh512-1 257-hue-HUE_SERVER]# cd /var/run/cloudera-scm-agent/process/`ls -valrt /var/run/cloudera-scm-agent/process | grep HUE_SERVER | tail -1 | awk '{print $9}'`
[root@hue4-cdh512-1 257-hue-HUE_SERVER]# pwd
/var/run/cloudera-scm-agent/process/257-hue-HUE_SERVER
[root@hue4-cdh512-1 257-hue-HUE_SERVER]# ll
total 76
-rwxr----- 1 hue hue 393 Aug 4 09:17 altscript.sh
-rw-r----- 1 hue hue 359 Aug 4 09:17 cloudera-monitor.properties
-rw----- 1 root root 21330 Aug 4 09:17 config.zip
-rw-r----- 1 hue hue 1987 Aug 4 09:17 creds.localjceks
drwxr-xr-x 2 hue hue 300 Aug 4 09:17 hive-conf
-rw-r----- 1 hue hue 4267 Aug 4 09:17 hue.ini
-rw-r----- 1 hue hue 99 Aug 4 09:17 hue.keytab
-rw-r----- 1 hue hue 0 Aug 4 09:17 hue_safety_valve.ini
-rw-r----- 1 hue hue 0 Aug 4 09:17 hue_safety_valve_server.ini
drwxr-x--x 2 hue hue 60 Aug 4 09:17 impala-conf
drwxr-x--x 2 hue hue 80 Aug 4 09:17 logs
-rw-r----- 1 hue hue 541 Aug 4 09:17 navigator.client.properties
-rw-r----- 1 hue hue 540 Aug 4 09:17 navigator.lineage.client.properties
-rw----- 1 root root 2415 Aug 4 09:17 proc.json
-rw-r----- 1 hue hue 0 Aug 4 09:17 redaction-rules.json
drwxr-x--x 2 hue hue 60 Aug 4 09:17 sentry-conf
-rw-r----- 1 hue hue 8328 Aug 4 09:17 service-metrics.properties
drwxr-x--x 2 hue hue 60 Aug 4 09:17 sqoop2-conf
-rw----- 1 root root 3000 Aug 4 09:17 supervisor.conf
drwxr-xr-x 2 hue hue 220 Aug 4 09:17 yarn-conf
```

The process directory for any given role is mirrored in Cloudera Manager. Go to Hue Instances , select a role such as Hue Server, and then click the Processes tab.

Related Information

[hue.ini](#)

Hue configurations in CDP Runtime

Cloudera Manager exposes properties that allow you to insert custom configuration text into XML configuration, property, and text files, or into an environment. The configuration snippet mechanism is intended for use in cases where there is a configuration setting that is not exposed as a configuration property in the Cloudera Manager web interface. Configuration snippets generally override normal configuration. You can configure custom properties for Hue by specifying them in the Advanced Configuration Snippet (Safety valve) field in Cloudera Manager.

The following table lists the safety valve parameters supported by Hue in CDP Runtime:

Parameter	Description
<code>[[desktop]] app_blacklist</code>	Used to add or remove applications, such as the File Browser, Impala, Hive, Oozie, and so on from the blocked list of applications.
<code>[desktop] [[session]] ttl=[**NUMBER-OF-SECONDS**]</code>	Used to configure the duration of a user session. The ttl property determines the length of time that the cookie with the user's session ID lives before expiring. After the ttl setting is reached, the user's session expires whether it is active or not. The default setting for ttl is 1,209,600 seconds, which equals two weeks.
<code>[jobbrowser] [[query_store]] server_url=[**QUERY-PROCESSOR-URL**]</code>	Used to display the Queries tab for Hive and Impala on the Job Browser page. This configuration is enabled by default and is not exposed in the Hue safety valve. However, to hide the Queries tab, you can override the configuration by adding the following lines in Cloudera Manager Clusters Hue Configuration Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini : <code>[jobbrowser] [[query_store]] is_enabled=false</code> To enable the Queries tab, set is_enabled to true.
<code>[aws] [[aws_accounts]] [[[default]]] access_key_id=[**AWS-ACCESS-KEY**] secret_access_key=[**SECRET-ACCESS-KEY**] region=[**AWS-REGION**]</code>	Used to enable the S3 File Browser for Hue without IDBroker.
<code>[aws] has_iam_detection=true [[aws_accounts]] [[[default]]] region=[**AWS-REGION**]</code>	Used to enable the S3 File Browser for Hue with IDBroker.
<code>[desktop] [azure] [[azure_accounts]] [[[default]]] client_id=[**AZURE-ACCOUNT-CLIENT-ID**] client_secret=[**AZURE-ACCOUNT-CLIENT-SECRET**] tenant_id=[**AZURE-ACCOUNT-TENANT-ID**] [[abfs_clusters]] [[[default]]] fs_defaultfs=abfs://[**CONTAINER-NAME**]@[**AZURE-STORAGE-ACCOUNT-NAME**]>.dfs.core.windows.net webhdfs_url=https://[**AZURE-STORAGE-ACCOUNT-NAME**].dfs.core.windows.net/</code>	Used to enable the ABFS File Browser for Hue without IDBroker.
<code>[desktop] [azure]</code>	Used to enable the ABFS File Browser for Hue with IDBroker.

Contact Cloudera Support if you are required to use a configuration snippet that is not explicitly documented.

Introduction to Hue Advanced Configuration Snippet (Safety valves)

To customize and configure Hue properties, such as authentication or enabling S3 or ABFS browsers, and so on that you cannot directly configure from the Cloudera Manager UI, you can use the Advanced Configuration Snippet field under the Hue cluster configurations in Cloudera Manager.

Advanced safety valves

The field for `hue_safety_valve.ini` is service-wide and affects all Hue roles listed on the Instances tab. The field for `hue_safety_valve_server.ini` only affects Hue Server role instances.

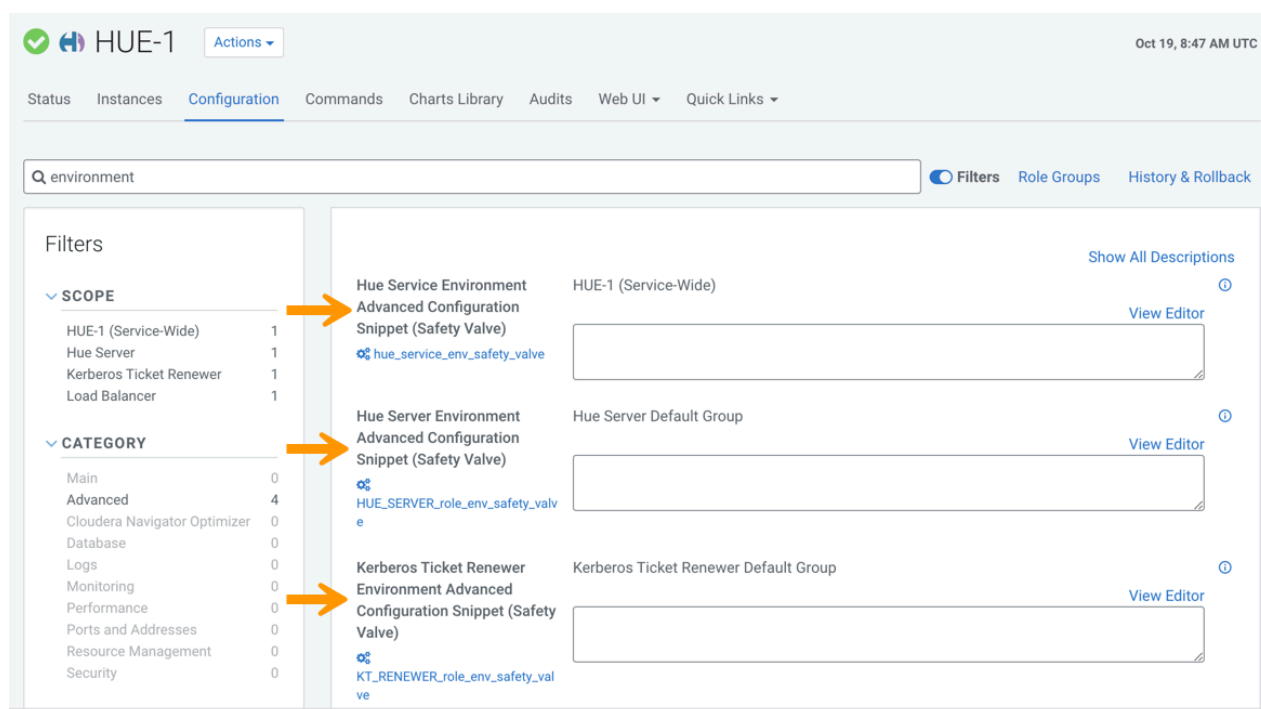
The screenshot shows the Cloudera Manager interface for the Hue-1 cluster. The 'Configuration' tab is active. In the 'Filters' sidebar, under 'SCOPE', 'HUE-1 (Service-Wide)' and 'Hue Server' are selected. The main content area displays two configuration snippets:

- Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini`**: This snippet is service-wide and affects all Hue roles. It contains the configuration for `hue_service_safety_valve`.
- Hue Server Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve_server.ini`**: This snippet only affects Hue Server role instances. It contains the configuration for `hue_server_hue_safety_valve`.

Arrows in the image indicate the navigation path: from the 'Configuration' tab to the search bar, then to the 'Hue Service' snippet, and finally to the 'Hue Server' snippet.

Environment safety valves

Environment safety-valves let you configure environment variables across the service or for specific role instances. These safety-valves lack headers.



About setting up the Hue SQL AI Assistant

Administrators are required to set up and enable the SQL AI Assistant before Analysts can use it to generate, edit, optimize, and fix queries using natural language in Hue.

First, you must obtain clearance from your organization's infosec team to make sure it is safe to use the SQL AI Assistant because some of the table metadata and data, as mentioned in the previous section, is shared with the LLM.

Next, prepare one of the following AI platforms of your choice and then configure the SQL AI Assistant in Hue:

- Microsoft Azure OpenAI service,
- Amazon Bedrock service, or
- OpenAI platform

Preparing the Microsoft Azure OpenAI service

Microsoft Azure allows dedicated deployments of OpenAI GPT models. Using Azure's OpenAI service is much more secure than the publicly hosted OpenAI APIs because the data can be processed in your Virtual Private Cloud (VPC) network. Due to security considerations, Cloudera recommends that you use GPT models in Hue SQL AI Assistant with Azure's OpenAI service.

Procedure

1. Obtain a Microsoft Azure subscription by working with your organization's IT team.
Subscriptions vary based on your team and purpose.
2. Register to access the Azure OpenAI service.
Azure OpenAI requires registration and is currently only available to approved enterprise customers and partners. Customers who wish to use Azure OpenAI are required to submit a [registration form](#).
3. Create an Azure OpenAI resource in the Azure portal.

4. Go to Overview from the left navigation and obtain the resource URL and resource keys from the Develop tab under the resource details page.
You can use any one of the two available keys.
5. Go to Azure OpenAI Studio at <https://oai.azure.com/portal> and create your deployment under Management Deployments .
6. Select gpt-35-turbo-16k or higher.

What to do next

Enable the SQL AI Assistant using Cloudera Manager.

Preparing the Amazon Bedrock Service

Amazon Bedrock is a fully managed service that makes foundation models from leading AI startups and Amazon available through an API.

Before you begin

You must have an AWS account with Bedrock access.

Procedure

1. Log in to the Amazon Bedrock service.
2. Obtain your access key and secret as follows:
 - a) Go to the IAM console: <https://console.aws.amazon.com/iam>.
 - b) Click on Users from the left menu and select the user you want to access.
 - c) Click on Security credentials.
 - d) Go to the Access keys section and note the access keys.
3. Establish Anthropic Claude access.

Claude from Anthropic is one of the best models available in Bedrock for SQL-related tasks. By default, Claude is not available on Bedrock. You need to place a special request for Claude.

After gaining access, you can try Claude in the text playground under the Amazon Bedrock service. If you are in the us-east-1 region, this must take you to <https://us-east-1.console.aws.amazon.com/bedrock/home?region=us-east-1#/text-playground>.

What to do next

Enable the SQL AI Assistant using Cloudera Manager.

Preparing the OpenAI platform

Learn how to set up SQL AI Assistant on the OpenAI platform.

Before you begin

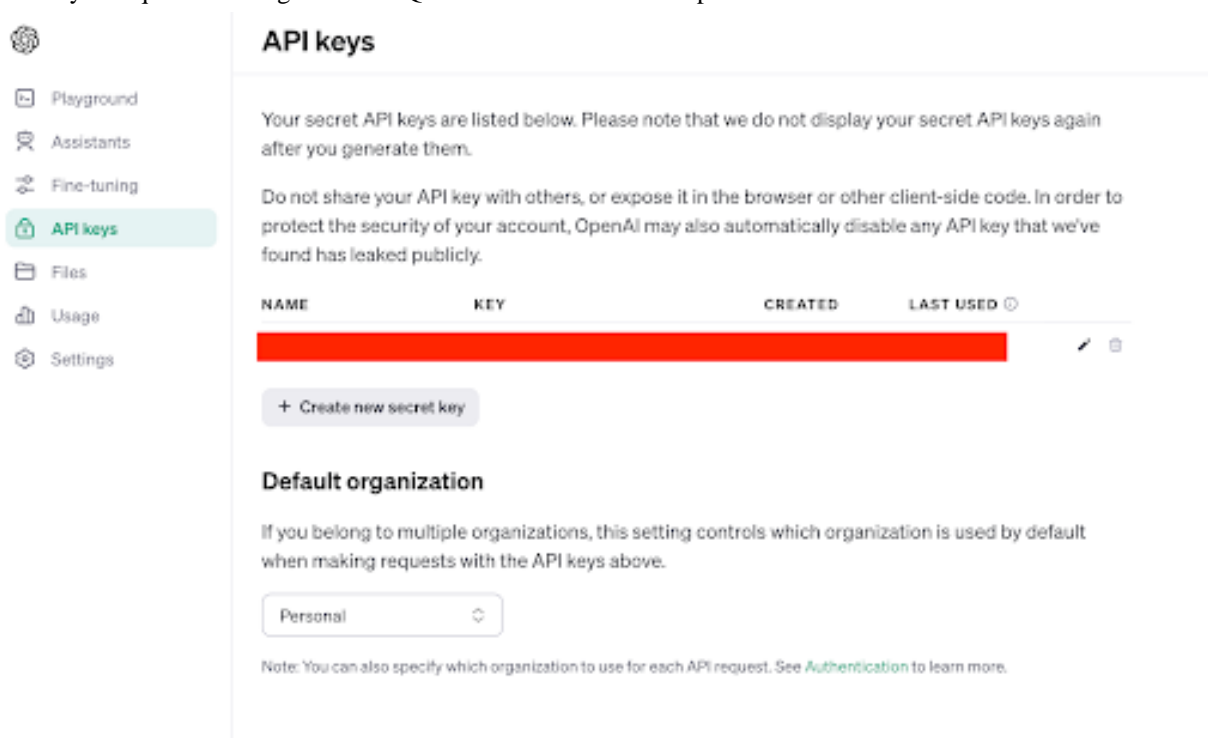
You must have created an account with the OpenAI platform.

Procedure

1. Log in to the OpenAI portal.

2. Obtain the API key by navigating to the API keys menu from the left navigation pane.

The key is required to integrate Hue SQL AI Assistant with the OpenAI service.



API keys

Your secret API keys are listed below. Please note that we do not display your secret API keys again after you generate them.

Do not share your API key with others, or expose it in the browser or other client-side code. In order to protect the security of your account, OpenAI may also automatically disable any API key that we've found has leaked publicly.

NAME	KEY	CREATED	LAST USED
[Redacted]			

+ Create new secret key

Default organization

If you belong to multiple organizations, this setting controls which organization is used by default when making requests with the API keys above.

Personal

Note: You can also specify which organization to use for each API request. See [Authentication](#) to learn more.

What to do next

Enable the SQL AI Assistant using Cloudera Manager.

Configuring the Hue SQL AI Assistant

Learn how to configure the Hue SQL AI Assistant.

About this task

In this task, you specify the token value in the Hue Advanced Configuration Snippet field in Cloudera Manager. The credentials are saved in a configuration file in the plain text format.



Note: Cloudera recommends that you use the open approach to pass tokens in test deployments, for proof of concept use cases. Use the secure approach in production deployments.

Procedure

1. Log in to the CDP Management Console as an Administrator.
2. Go to Environments and select your environment.
3. Go to the Data Lake tab and click on the CM URL to open Cloudera Manager.
4. Go to Clusters Hue service Configuration and select add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini field:

Azure

```
[desktop]
[[ai_interface]]
model_name='[***DEPLOYMENT-NAME***]'
base_url="https://[***RESOURCE***].openai.azure.com/"
```

```
token=" [***RESOURCE-KEY*** ]"
```

AWS

```
[aws]
[[bedrock_account]]
  access_key_id=' [***ACCESS-KEY*** ]'
  secret_access_key=' [***SECRET-KEY*** ]'
  region='us-east-1'
[desktop]
[[ai_interface]]
  service='bedrock'
  model='claude'
```



Note: You can also use a newer model such as Claude 3 by specifying model='claude3'.

OpenAI

```
[desktop]
[[ai_interface]]
  service='openai'
  token=' [***API-KEY*** ]'
```

5. Click Save Changes.

Results

You see ✨ Assistant on the Hue SQL editor.

Service and model-related configurations for setting up the Hue SQL AI Assistant

Review the list of service, model, and semantic search-related configurations used for custom configuring the AI services and models you want to use with the SQL AI Assistant and how to specify them in the Hue Advanced Configuration Snippet in Cloudera Manager.

List of service and model-related configurations

You can configure the AI services and models you want to use by going to Cloudera Manager Clusters Hue service Configurations Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini and adding the following lines:

```
[desktop]
[[ai_interface]]
  [***CONFIG-KEY1***]=' [***VALUE*** ]'
  [***CONFIG-KEY2***]=' [***VALUE*** ]'
[[semantic_search]]
  [***CONFIG-KEY1***]=' [***VALUE*** ]'
  [***CONFIG-KEY2***]=' [***VALUE*** ]'
```

Specify the service and model-related configurations under the [[ai_interface]] section as listed in the following table:

AI interface config key	Description
service	API service to be used for AI tasks. AI is disabled when a service is not configured. For example, ai_service and ai_assistant.
trusted_service	Indicates whether the LLM is trusted or not. Turn on to disable the warning. The default value is false.

AI interface config key	Description
model	The AI model you want to use for AI tasks. For example, gptand llama.
model_name	The fully qualified name of the model to be used. For example, gpt-3.5-turbo-16k.
base_url	Service API base URL.
add_table_data	When enabled, sample rows from the table are added to the prompt. The default value is True.
table_data_cache_size	Size of the LRU cache used for storing table sample data.
auto_fetch_table_meta_limit	Number of tables to load from a database, initially.
token	Service API secret token.
token_script	Provides a secure way to get the service API secret token.

List of semantic search-related configurations

Specify the semantic search-related configurations used for RAG under the [[semantic_search]] section, as listed in the following table:

Semantic search config key	Description
relevancy	The technology you want to use for semantic search. Acceptable values are vector_search or v
embedding_model	The model you want to use for data-embedding. This must be compatible with SentenceTransf
cache_size	Size of the LRU cache used for storing embedding.

Hue logs

Cloudera Manager generates standard stream logs when each Hue role instance starts and stops. The Hue service, which is built on the [Django framework](#), generates log4j logs when the service is running.

Standard stream logs

Cloudera Manager logs the start and stop of each supervised Hue process in standard stream logs (stdout.log, stderr.log).

When the Hue service restarts, Cloudera Manager generates a new directory for each supervised process of a Hue role instance:

```
ls -vrl /var/run/cloudera-scm-agent/process | grep HUE
```

```
[root@hue4-cdh512-1 ~]# ls -vrl /var/run/cloudera-scm-agent/process | grep HUE
drwxr-x--x 4 hue hue 320 Aug 6 14:33 289-hue-HUE_LOAD_BALANCER
drwxr-x--x 8 hue hue 440 Aug 6 14:33 288-hue-HUE_SERVER
drwxr-x--x 4 hue hue 280 Aug 6 14:33 258-hue-HUE_LOAD_BALANCER
drwxr-x--x 8 hue hue 460 Aug 6 14:33 257-hue-HUE_SERVER
drwxr-x--x 4 hue hue 280 Aug 4 09:17 242-hue-HUE_LOAD_BALANCER
drwxr-x--x 8 hue hue 460 Aug 4 09:17 241-hue-HUE_SERVER
drwxr-x--x 4 hue hue 280 Aug 4 08:41 239-hue-HUE_LOAD_BALANCER
drwxr-x--x 8 hue hue 460 Aug 4 08:41 238-hue-HUE_SERVER
drwxr-x--x 4 hue hue 280 Aug 4 08:38 236-hue-HUE_LOAD_BALANCER
drwxr-x--x 8 hue hue 460 Aug 4 08:38 235-hue-HUE_SERVER
```

It writes to a nested logs directory for each running instance:

```
[root@hue4-cdh512-1 logs]# pwd
/var/run/cloudera-scm-agent/process/289-hue-HUE_LOAD_BALANCER/logs
[root@hue4-cdh512-1 logs]# ll
total 16
-rw-r--r-- 1 root root 11148 Aug  6 14:33 stderr.log
-rw-r--r-- 1 root root  447 Aug  6 14:33 stdout.log
```

Configuration errors are written here because they prevent Hue servers and load balancers from starting properly.



Tip: Testing the LDAP configuration from Cloudera Manager (Clusters Hue service Test LDAP Configuration) also writes to standard stream logs which you can search using the following command: `ls -vrl /var/run/cloudera-scm-agent/process | grep ldaptest`

Hue service Django logs

When the Hue service is running, Hue generates logs in `/var/log/hue` using log4j. Load balancer logs are in `/var/run/httpd`. You can view these logs in Hue at `http://hueserver:port/logs`.

Table 1: Hue service logs

Log Name	Description
access.log	Filtered list of successful attempts to access Hue Web UI.
audit.log	Captures user operations, such as: <ul style="list-style-type: none"> • USER_LOGIN • USER_LOGOUT • EXPORT • DOWNLOAD • DELETE_USER • DELETE_GROUP • EDIT_USER • CREATE_USER • EDIT_GROUP • CREATE_GROUP • EDIT_PERMISSION • ADD_LDAP_USERS • ADD_LDAP_GROUPS • SYNC_LDAP_USERS_GROUPS
audit/hue_server_audit_wal.log	Audit log visible in Apache Atlas.
error.log	Filtered list of all nontrivial errors
kt_renewer.log	Kerberos ticket renews
metrics-hue_server/metrics.log	Populates charts in Cloudera Manager
migrate.log	Database and table migrations + First Run of Hue server
rungunicornserver.log	Hue (Gunicorn) web server information.
hue_install.log	Contains the log produced during installation

Enabling DEBUG logging for Hue logs

DEBUG is available for the Hue server logs in `/var/log/hue`. By default, the Hue service writes INFO level messages and keeps a small buffer of log messages at all levels in memory. If you want detailed, debug-level information, then you can enable this feature from Hue configuration in Cloudera Manager.

About this task

In the debug mode, Hue displays a detailed traceback when an exception occurs. Debugging information may contain sensitive data.

Procedure

1. Log in to the CDP web interface as an Administrator.
2. Go to **Management Console Environments**, select your environment and go to the **Data Lake** tab.
3. Click on the CM URL to open Cloudera Manager.
4. Go to **Clusters Hue Configuration** and select the **Enable Django Debug Mode** option.
The **Enable Django Debug Mode** option is used to toggle between INFO and DEBUG-level logging.
5. Restart the Hue service.

Enabling httpd log rotation for Hue

Several log files, such as access logs and error logs, are produced by the Apache HTTP server. Over time, these files can grow in size and can cause disk issues. You can prevent this problem by using Apache's `rotatelogs` utility and defining when to rotate logs – either based on a time interval or maximum size of the log.

About this task

You can configure the Apache server to rotate the logs when a certain filesize is reached (for example, 5 Bytes or 10 KBytes) or at a set time interval (for example, daily or weekly), or both. If you want to specify both the time and size, the size must be specified after the time. Log rotation occurs whenever time or size limits are reached.

The log files are generated as per the configurations defined in the `httpd.conf` file.

Before you begin

The `rotatelogs` utility must be present on your Operating System, and you must know its location so that you can use it. Run the following command from the Hue load balancer host system to get the path of the `rotatelogs` utility:

```
[root@example-domain-hue-HUE_LOAD_BALANCER]# which rotatelogs
/usr/sbin/rotatelogs
```

Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to **Clusters Hue service Configuration**.
3. Configure the log rotation policy based on your criteria in the **Load Balancer Advanced Configuration Snippet (Safety Valve)** for `httpd.conf` field.

To rotate logs whenever the specific file size is reached, add the following lines:

```
CustomLog " | /usr/sbin/rotatelogs /var/log/hue-httpd/access_log.%Y-%m-%d-%H_%M_%S [***FILE-SIZE***] (B|K|M|G)" common
```

```
ErrorLog "|/usr/sbin/rotatelog /var/log/hue-httpd/error_log.%Y-%m-%d-%H_%M_%S [***FILE-SIZE***](B|K|M|G)"
```

For example:

```
CustomLog "|/usr/sbin/rotatelog /var/log/hue-httpd/access_log.%Y-%m-%d-%H_%M_%S 5B" common
ErrorLog "|/usr/sbin/rotatelog /var/log/hue-httpd/error_log.%Y-%m-%d-%H_%M_%S 5B"
```

This configuration rotates the log file whenever it reaches a size of 5 bytes, and the suffix to the log file name will be created with the format access_log.YYYY-mm-dd-HH_MM_SS and error_log.YYYY-mm-dd-HH_MM_SS.

To rotate logs daily, add the following lines:

```
CustomLog "|/usr/sbin/rotatelog -l /var/log/hue-httpd/access_log.%Y-%m-%d-%H_%M_%S 86400" common
ErrorLog "|/usr/sbin/rotatelog -l /var/log/hue-httpd/error_log.%Y-%m-%d-%H_%M_%S 86400"
```

This creates the files /var/log/access_log.yyyy.mm.dd and /var/log/error_log.yyyy.mm.dd where yyyy is the year, mm is the month, and dd is the day of the month. The log operation switches to a new file every day at midnight, local time.

For more information, see *rotatelog - Piped logging program to rotate Apache logs*.

4. Click Save Changes.
5. Restart the affected services.

Related Information

[rotatelog - Piped logging program to rotate Apache logs](#)

Hue supported browsers

Hue works with the two most recent [LTS](#) (long term support) or [ESR](#) (extended support release) browsers. Cookies and JavaScript must be enabled.

The lists the minimum tested versions of the most common browsers:

- Chrome: ([Version history](#))
- Firefox: ([Version history](#))
- Safari (Mac only): [Version history](#)
- Microsoft Edge: ([Version history](#))

Hue can display in other browsers and in older versions of the common browsers, but you might not have access to all features.

Enabling cache-control HTTP headers when using Hue

You can enable Hue to use HTTP headers such as Cache-Control, Pragma, and Expires to ensure that your browser always uses the fresh version of the resource. You must set the value of the custom_cache_control property to true in Hue's Advanced Configuration Snippet.

About this task



Note: Enabling the use of these cache-control HTTP headers can affect performance because your browser tries to fetch the latest resources instead of using cached data. To disable cache-control HTTP headers, set the value of the `custom_cache_control` property to `false`.

Procedure

1. Log in to the CDP web interface as an Administrator.
2. Go to Management Console Environments , select your environment and go to the Data Lake tab.
3. Click on the CM URL to open Cloudera Manager.
4. Go to Clusters Hue Configuration and add the following line in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[desktop]
custom_cache_control=true
```

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

Setting up a Hue service account with a custom name

In Cloudera, “hue” username is the default machine user and group name for Hue. You can create a custom machine user and group name on the Hue cluster and then specify them in the Hue configurations in Cloudera Manager.

Procedure

1. SSH into the Hue server host as a user with sudo privileges
2. Create a new group for your custom Hue user account by running the following command:

```
sudo groupadd [***NEW-HUE-GROUP-NAME***]
```

Replace `[***NEW-HUE-GROUP-NAME***]` with the custom group name you want to use.

3. Create a new user and assign it to the newly created group by running the following command:

```
sudo useradd -G [***NEW-HUE-GROUP-NAME***] -m [***NEW-HUE-USER-NAME***] -s /usr/sbin/nologin
```

The new user is stored in the `/etc/passwd` file, and group is stored in the `/etc/group` file.

4. Verify that the new user exists within the newly created group by running the following command:

```
grep [***NEW-HUE-USER-NAME***] /etc/passwd
```

5. Change the ownership of the Hue directories as follows:

```
sudo chown -R [***NEW-HUE-USER-NAME***]:[***NEW-HUE-GROUP-NAME***] /var/log/hue
sudo chown -R [***NEW-HUE-USER-NAME***]:[***NEW-HUE-GROUP-NAME***] /var/log/hue-httpd
sudo chown -R [***NEW-HUE-USER-NAME***]:[***NEW-HUE-GROUP-NAME***] /tmp/hue_*
```

6. Modify the Hue configuration in Cloudera Manager as follows:

- a) Log in to Cloudera Manager as an Administrator.
- b) Go to Clusters Hue Configuration .
- c) Specify the custom group name you created earlier[***NEW-HUE-GROUP-NAME***] in the Default User Group field.
- d) Specify the custom user name you created earlier[***NEW-HUE-USER-NAME***] in the System User field.
- e) Specify the custom group name you created earlier[***NEW-HUE-GROUP-NAME***] in the System Group field.
- f) Add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini field.

```
[desktop]
server_user=[ ***NEW-HUE-USER-NAME*** ]
server_group=[ ***NEW-HUE-GROUP-NAME*** ]
default_user=[ ***NEW-HUE-USER-NAME*** ]
```

- g) Click Save Changes.
- h) Restart the Hue service.

7. Tail the /var/log/hue/rungunicornserver.log log file for any issues.**8. Change the permissions of the Hue home directory as follows:**

```
sudo cd /opt/cloudera/parcels/CDH/lib/
sudo chown -R [ ***NEW-HUE-USER-NAME*** ]:[ ***NEW-HUE-GROUP-NAME*** ]hue
```

9. Change the permissions of the process directory as follows:

```
sudo cd /var/run/cloudera-scm-agent/
sudo setfacl -Rm user:[ ***NEW-HUE-USER-NAME*** ]:rwx process
```

10. Verify Access Control List (ACL) changes by running the following command:

```
getfacl process
```

To validate newly configured custom user and group

Run the following Python script on your Hue hosts to validate that the new user and group names are valid and configured correctly.

This script requires you to input the user (uid) and group (gid) IDs corresponding to the custom user and group names you created and displays the effective user and group IDs that are in use. You must ensure that the user belongs to the same group while specifying the uid and gid. The script displays the “Permission denied: Operation not permitted” error if there is a mismatch between the user ID and the group ID :

```
import os
def change_user_and_group(euid, egid):
    try:
        # Change the effective group ID
        os.setegid(egid)
        print(f"Effective Group ID changed to: {egid}")
        # Change the effective user ID
        os.seteuid(euid)
        print(f"Effective User ID changed to: {euid}")
        # Verify changes
        print(f"Current Effective User ID: {os.geteuid()}")
        print(f"Current Effective Group ID: {os.getegid()}")
    except PermissionError as e:
        print(f"Permission denied: {e}")
```

```

except Exception as e:
    print(f"An error occurred: {e}")

if __name__ == "__main__":
    # Replace with actual user and group IDs that you want to switch to
    # Change `new_euid` to the user ID (UID) of the custom user you created.
    # Change `new_egid` to the group ID (GID) of the custom group you create
    d.
    # Example: new_euid = 1001 # Non-root user ID
    #           new_egid = 1001 # Non-root group ID
    new_euid = 1001 # Replace with the new user ID (UID)
    new_egid = 1001 # Replace with the new group ID (GID)
    change_user_and_group(new_euid, new_egid)

```

Options to restart the Hue service

Restarting the Hue service is necessary after upgrading CDP or making configuration changes to the Hue service. You can restart Hue in regular or rolling restart mode. Rolling restart reduces the service downtime to approximately 80 to 90 seconds.


You can restart the Hue service independently or as a part of the cluster restart. When you restart the Hue service normally, Hue can become unavailable for an average of 30 minutes from the restart time (this could vary depending on the number of services in your cluster). When you restart the cluster containing the Hue service, Hue is stopped first among other services in the cluster and restarted at the end. To reduce the service downtime, Cloudera recommends that you use the “Rolling Restart” option.



Note: Cluster-level rolling restart is supported only in High Availability (HA) clusters. However, you can restart the Hue service and its roles independently in the rolling restart mode in both HA and non-HA clusters.

Restarting the Hue service independently in rolling restart mode

To restart Hue service in the rolling restart mode, use one of the following steps:


- Go to **Cloudera Manager Clusters** and click  corresponding to the Hue service and then click **Rolling Restart**, or
- Go to **Cloudera Manager Clusters Hue** and click **Actions Rolling Restart**.

Cloudera Manager shows the following options to restart the Hue service in the rolling restart mode:

Category	Restart options	Description
Restart filters	Restart roles with stale configuration only	Select this option to restart Hue if you have updated the Hue configuration.
	Restart roles with the old software version only	Select this option to restart Hue if you have upgraded the cluster.
Role types to restart	Kerberos Ticket Renewer	Selected by default. Restarts the Kerberos Ticket Renewer role.
	Load Balancer	Selected by default. Restarts the Load Balancer role after the Kerberos Ticket Renewer.
	Hue Server	Selected by default. Restarts the Hue Server role after the load balancer.

When you restart Hue, Hue’s non-worker roles, such as the load balancer, Kerberos ticket renewer, and Hue server restart one after the other.

Restarting a CDP cluster containing the Hue service in the rolling restart mode

To restart an HA cluster containing the Hue service in the rolling restart mode, go to **Cloudera Manager Clusters** and click  corresponding to your cluster name and select **Rolling Restart**. You see the **Rolling Restart** modal with the list of services that support rolling restart.



Note: All Hue roles (namely the load balancer, Kerberos ticket renewer, and Hue server) are of the non-worker type. You must select one of the following options under the **Roles** to include section for ensuring that the Hue service is restarted in the rolling restart mode, depending on your requirements:

- Non-Workers Only, or
- All Roles

Customizing the Hue web interface

To customize the Hue web interface, add configuration properties in Cloudera Manager. You can customize the banner, the page logo, the cache timeout setting, and you can enable or disable anonymous usage data collection.

Adding a custom banner in Hue

You can add a custom banner to the Hue web interface by adding your custom HTML to the **Top Banner Custom HTML** property in Cloudera Manager.

Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to **Clusters Hue service Configuration** and add the banner text in the **Top Banner Custom HTML** field as follows:
`<H1>Your company's custom Hue Web UI banner</H1>`
3. Click **Save Changes**.
4. Restart the Hue service.

Changing the page logo in Hue

You can replace the Hue web interface logo with a custom log that is created with an SVG code. Add any type of logo you want, but your custom logo should be designed to fit into a 160 x 40 pixel space.

About this task

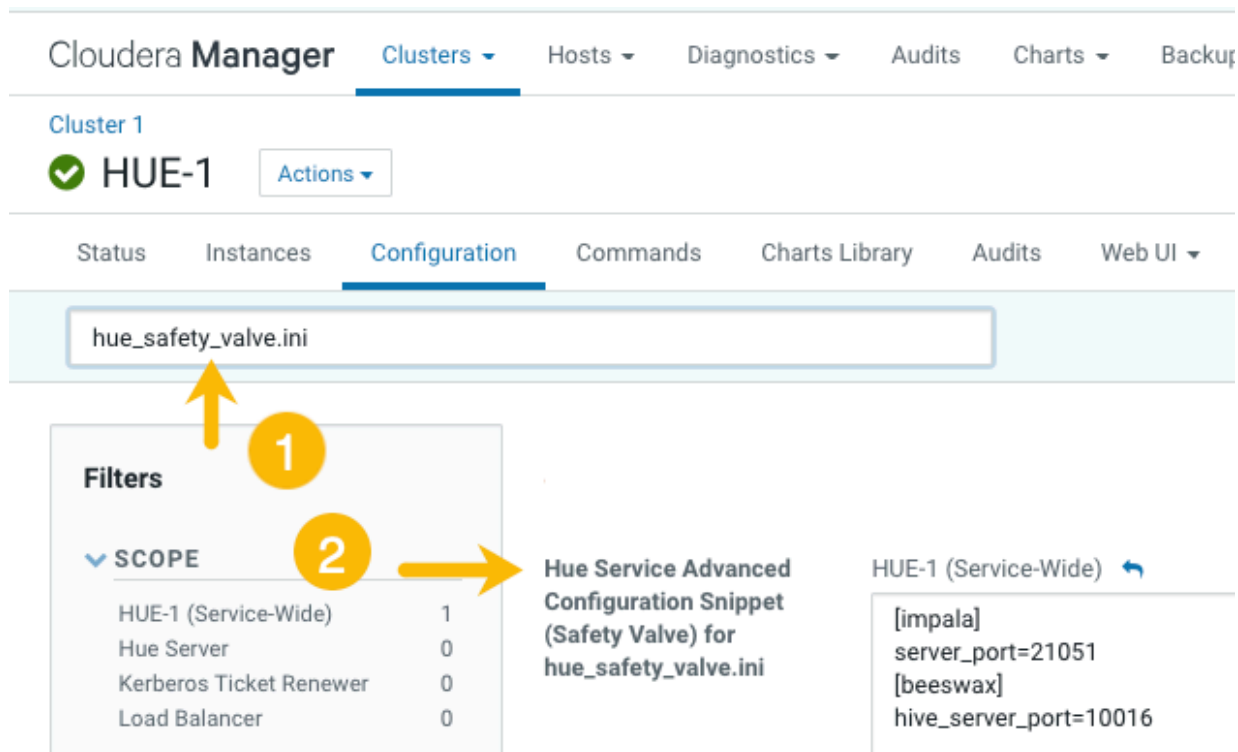
For example, here is the Hue logo shown in the following image:



You can change this Hue logo by adding the appropriate SVG code to the `logo_svg` property under the `[desktop]` `[[custom]]` section in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` configuration parameter in Cloudera Manager.

Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue Configuration
3. In the Search text box, type `hue_safety_valve.ini` to locate the configuration parameter:



4. Add the following parameters with your custom logo SVG code to the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini configuration parameter:

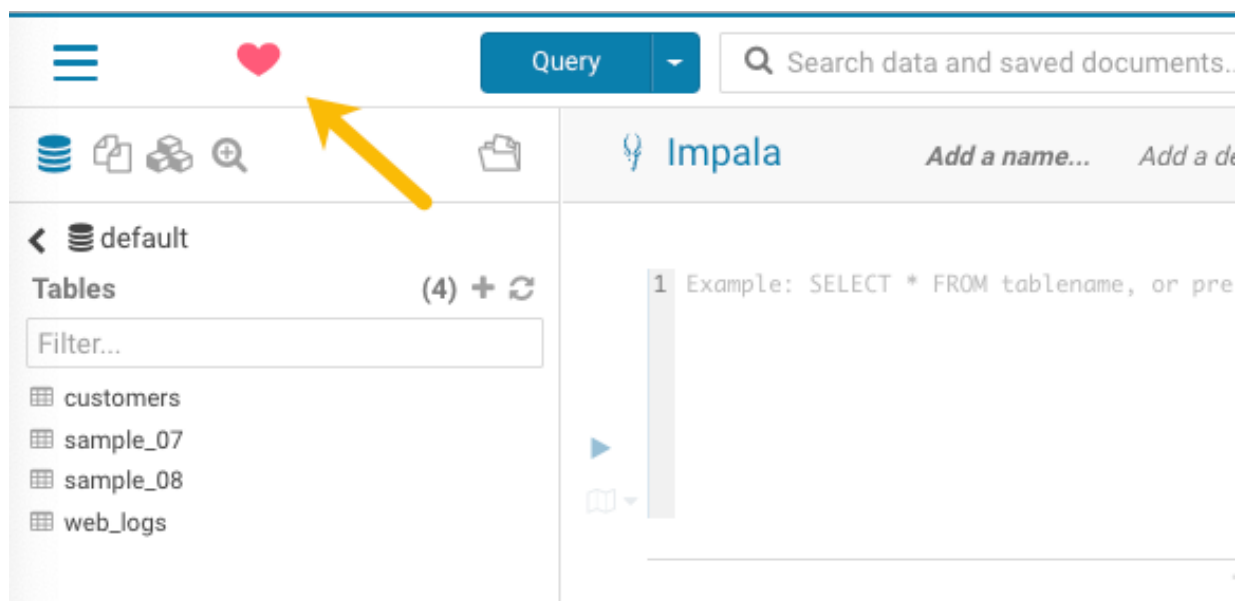
```
[desktop]
[[custom]]
logo_svg=' [ ***SVG-CODE-FOR-CUSTOM-LOGO*** ] '
```

For example, the following SVG code replaces the Hue logo with a red heart:

```
[desktop]
[[custom]]
logo_svg='<g><path stroke="null" id="svg_1" d="m44.41215,11.43463c-4.05017,-10.71473 -17.19753,-5.90773 -18.41353,-0.5567c-1.672,-5.70253 -14.497,-9.95663 -18.411,0.5643c-4.35797,11.71793 16.891,22.23443 18.41163,23.95773c1.5181,-1.36927 22.7696,-12.43803 18.4129,-23.96533z" fill="#ffffff"/> <path stroke="null" id="svg_2" d="m98.41246,10.43463c-4.05016,-10.71473 -17.19753,-5.90773 -18.41353,-0.5567c-1.672,-5.70253 -14.497,-9.95663 -18.411,0.5643c-4.35796,11.71793 16.891,22.23443 18.41164,23.95773c1.5181,-1.36927 22.76959,-12.43803 18.41289,-23.96533z" fill="#FF5A79"/> <path stroke="null" id="svg_3" d="m154.41215,11.43463c-4.05016,-10.71473 -17.19753,-5.90773 -18.41353,-0.5567c-1.672,-5.70253 -14.497,-9.95663 -18.411,0.5643c-4.35796,11.71793 16.891,22.23443 18.41164,23.95773c1.5181,-1.36927 22.76959,-12.43803 18.41289,-23.96533z" fill="#ffffff"/> </g>'
```

5. Click Save Changes.
6. Restart the Hue service.
7. In the Hue configuration page of Cloudera Manager, select **Web UI Hue Load Balanced** to load Hue and view your custom logo.

If you added the sample SVG code that defines a red heart as the logo, your Hue web UI looks like this:



- 8.
- 9.
- 10.

Related Information

[Scalable Vector Graphics](#)

Adding a splash screen in Hue

You can add a custom splash screen to the Hue web interface by adding your custom HTML to the Advanced Configuration Snippet in Cloudera Manager.

Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to **Clusters Hue service Configuration** and add the banner test in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini field as follows:

```
[desktop]
[[custom]]
login_splash_html=[ ***CUSTOM-HTML*** ]
```

```
[desktop]
[[custom]]
login_splash_html=<h1>Hue, the next-gen SQL Assistant</h1>
```

3. Click **Save Changes**.
4. Restart the Hue service.

Setting the cache timeout

Enable Hue UI caching by setting a timeout value in milliseconds. The default is 10 days or 864000000 milliseconds. Set the timeout to 0 to disable caching.

About this task

When you browse tables using the left assist panel or run queries, Hue caches this information for fetching information faster and query autocompletion. You can configure the time for which you want to Hue to cache this information by setting the value of the cacheable_ttl property under the [desktop][[custom]] section in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini configuration property in Cloudera Manager.

Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to **Clusters Hue Configuration**.
3. In the Search text box, type hue_safety_valve.ini to locate the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini configuration parameter.
4. Add the following parameters with the cache timeout value to the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini configuration parameter:

```
[desktop]
[[custom]]
cacheable_ttl=[ ***VALUE-IN-MILLISECONDS*** ]
```

For example, the following configuration sets the cache timeout to the default value of 86400000 milliseconds:

```
[desktop]
[[custom]]
```

```
cacheable_ttl=86400000
```

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

Enabling or disabling anonymous usage data collection

Hue tracks anonymized pages and application versions to gather information about application usage levels. The data collected does not include hostnames or IDs. For example, the data collected has the format /2.3.0/pig or /2.5.0/beeswax/execute.

About this task

To enable or disable anonymous usage data collection:

Procedure

1. In the Cloudera Manager Admin Console, select `ClustersHueConfiguration` to navigate to the configuration page for Hue.
2. In the Search text box, type `usage` to locate the Enable Usage Data Collection check box:
 - To enable anonymous data collection, check the box, which is the default setting.
 - To disable anonymous data collection, clear the check box.
3. Enter a Reason for change..., and then click Save Changes at the bottom of the page to save the configuration change.
- 4.



Refresh the browser page and click the restart icon at the top of the page so the new configuration changes can be read by the server and the new data collection setting takes effect.

Disabling the share option in Hue

Hue allows you to share documents, queries, and workflows with other users, either for viewing only or viewing and modifying. Sharing is enabled by default in the Hue UI. For added privacy and control, you can disable sharing by setting the `enable_sharing` property to `false` in the Advanced Configuration Snippet.

About this task

The sharing option is always available to the admin users. To disable the share option:

Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to `Clusters Hue service Configuration Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini` and set the value of the `enable_sharing` property to `false` in the `[desktop]` section:

```
[desktop]
enable_sharing=false
```

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

Enabling Hue applications with Cloudera Manager

Most Hue applications are configured by default, based on the services you have installed. Cloudera Manager selects the service instance that Hue depends on. If you have more than one service, you may want to verify or change the service dependency for Hue. If you add a service such as Oozie after you have set up Hue, you must set the dependency because it is not done automatically.

About this task

To add a dependency in Hue:

Procedure

1. In the Cloudera Manager Admin Console, select `ClustersHueConfiguration` to navigate to the configuration page for Hue.
2. Filter by `ScopeHue (Service-Wide)` and `CategoryMain`.
3. Select the `<service_name>` Service property that you want to set a dependency for. Select none to remove a dependency.
4. Enter a Reason for change..., and then click `Save Changes` at the bottom of the page to save the configuration change.
- 5.



Refresh the browser page and click the restart icon at the top of the page so the new configuration changes can be read by the server and the new setting takes effect.

Running shell commands

You can run shell commands to administer Hue programmatically. For example, to reset the superuser password or to assign an LDAP user superuser permissions.

About this task

To run Hue shell commands:

Procedure

1. Set `HUE_CONF_DIR` to the latest Hue process directory:

```
export HUE_CONF_DIR="/var/run/cloudera-scm-agent/process/`ls -alrt /var/run/cloudera-scm-agent/process | grep HUE_SERVER | tail -1 | awk '{print $9}'`"
echo $HUE_CONF_DIR
```

2. Set environment variables used to run the Hue webserver:

- CentOS/RHEL:

```
for line in `strings /proc/$(lsof -i :8888|grep -ml python|awk '{ print $2 }')/environ|egrep -v "^HOME=|^TERM=|^PWD=|^CLDR_JAVA_OPTS="`;do export $line;done
```

- Ubuntu:

```
for line in `strings /proc/$(lsof -i :8888|grep -ml hue|awk '{ print $2 }')/environ|egrep -v "^HOME=|^TERM=|^PWD=|^CLDR_JAVA_OPTS="`;do export $line;done
```



Note: Environment variables separated by spaces or underscores (_) are not parsed correctly, and must be excluded by using a caret symbol (^) to prevent errors.

3. Run shell subcommands

When true, HUE_IGNORE_PASSWORD_SCRIPT_ERRORS runs the Hue shell even if hue.ini contains passwords generated by Cloudera Manager (such as bind_password and ssl_password).



Note: Do not export HUE_IGNORE_PASSWORD_SCRIPT_ERRORS or HUE_DATABASE_PASSWORD to ensure that they are not stored and only apply to this command.

For CDH parcel deployments:

- Run the interactive Hue Python shell (Ctrl+D to quit)

```
HUE_IGNORE_PASSWORD_SCRIPT_ERRORS=1 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/hue shell
```

Or with the database password:

```
HUE_IGNORE_PASSWORD_SCRIPT_ERRORS=1 HUE_DATABASE_PASSWORD=<your db password> /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/hue shell
```

- Change a user password

```
HUE_IGNORE_PASSWORD_SCRIPT_ERRORS=1 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/hue changepassword admin
```

- Promote Hue user to superuser

```
HUE_IGNORE_PASSWORD_SCRIPT_ERRORS=1 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/hue shell
```

```
from django.contrib.auth.models import User
a = User.objects.get(username='gwen')
a.is_superuser = True
a.save()
```

- Count all of the documents of a certain user:

```
from django.contrib.auth.models import User
from desktop.models import Document2

user=User.objects.get(username='demo')
Document2.objects.documents(user=user).count()

Out[X]: 1167
```

- List available subcommands

```
HUE_IGNORE_PASSWORD_SCRIPT_ERRORS=1 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/hue
```

For CDH package deployments:

- ```
HUE_IGNORE_PASSWORD_SCRIPT_ERRORS=1 /usr/lib/hue/build/env/bin/hue shell
```

## Downloading and exporting data from Hue

Hue enables you to download or export data from Hue to HDFS or to an external storage location from Hue Editor, Hue Dashboard, and the Hue File browser. You can limit the number of rows or bytes that are downloaded or disable the export feature altogether so that you do not run out of storage space.

For a service-wide change, go to [Cloudera Manager Clusters Hue service Configuration](#) and specify the configurations in the Hue Service Advanced Configuration Snippet (Safety valve) for `hue_safety_valve.ini` field.

By default, Hue users can download the query results from the Hue Editor, the Hue Dashboard, and the File browser.

### Limiting the number of rows to download

Specify the following in the Hue Service Advanced Configuration Snippet (Safety valve) for `hue_safety_valve.ini` to limit the number of rows that can be downloaded from a query before it is truncated:

```
[beeswax]
download_row_limit=x
```

`x` represents the number of rows that you can download.

By default, there is no download limit, and you can configure this by setting the value to “-1”:

```
[beeswax]
download_row_limit=-1
```

### Limiting the number of bytes to download

Specify the following in the Hue Service Advanced Configuration Snippet (Safety valve) for `hue_safety_valve.ini` to limit the number of bytes that can be downloaded from a query before it is truncated:

```
[beeswax]
download_bytes_limit=x
```

`x` represents the number of bytes that you can download.

By default, there is no download limit, and you can configure this by setting the value to “-1”:

```
[beeswax]
download_bytes_limit=-1
```

### Disabling the data download feature

Specify the following in the Hue Service Advanced Configuration Snippet (Safety valve) for hue\_safety\_valve.ini field to disable your users from downloading query results:

```
[desktop]
enable_download=false
```

Specify the following in the Hue Service Advanced Configuration Snippet (Safety valve) for hue\_safety\_valve.ini field to hide the Download button from the Hue File browser:

```
[filebrowser]
show_download_button=false
```

## Enabling a multi-threaded environment for Hue

A multi-threaded environment can help reduce the time it takes to migrate data from one database to other. By default, operations such as migrating data run on a single thread. For example, if you are switching from MySQL as the backend database for Hue to Oracle, then enabling a multi-threaded environment significantly reduces the data transfer time.

### Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini .
3. Locate the [[database]] section under [desktop] and set threaded to true:

```
[desktop]
[[database]]
options={"threaded":true}
```

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

## Adding Query Processor service to a cluster

The Query Processor service indexes Hive and Tez events and provides APIs to access them. It is required if you want to view the Queries tab (query history and query details) on the Hue Job Browser.

### About this task



**Attention:** This task is applicable only if you are upgrading to CDP 7.1.8 and higher. If you are upgrading to CDP 7.1.7 SP2, 7.1.7 SP1, 7.1.7, or lower, then you can skip this task.

You can either install the Query Processor service as a dependency while adding the Hue service or after adding the Hue service.

### Before you begin



**Note:** You must be assigned the DataHubCreator role to perform this task.

This task assumes that you already have a database installed on a host in your cluster. On CDP, Hue Query Processor supports PostgreSQL, MySQL, and Oracle databases. The supported PostgreSQL database version for Hue Query Processor is 9.6 and higher.

Next, you need to create a database for the Query Processor service with the required roles. To create the Query Processor database:

1. SSH into your database host as a root user.
2. Start the database terminal by running the following commands:

(PostgreSQL)

```
sudo -u postgres psql
```

(MySQL)

```
systemctl start mysqld
```

(Oracle)

```
sqlplus / as sysdba
```

3. Run the following statement while specifying the username, password, and a database name for the Query Processor:

(PostgreSQL)

```
CREATE ROLE [***QP-USER***] WITH LOGIN PASSWORD '[***QP-PASSWORD***]';
ALTER ROLE [***QP-USER***] CREATEDB;
CREATE DATABASE [***QP-DATABASE***];
GRANT ALL PRIVILEGES ON DATABASE [***QP-DATABASE***] TO [***QP-USER***];
```

(MySQL)

```
CREATE USER [***QP-USER***] IDENTIFIED BY '[***QP-PASSWORD***]';
CREATE DATABASE [***QP-DATABASE***];
USE [***QP-DATABASE***];
GRANT ALL PRIVILEGES ON [***QP-DATABASE***].* TO [***QP-USER***];
```

(Oracle)

```
CREATE USER [***QP-USER***] IDENTIFIED BY [***QP-PASSWORD***];
GRANT ALL PRIVILEGES TO [***QP-USER***];
```



**Note:** When you create a user in the Oracle database, the associated schema is created automatically. This schema becomes the default schema for that user.

## Procedure

1. Log in to the CDP Management Console as an Administrator.
2. Go to your environment and click on CM UI to open Cloudera Manager.
3. Go to Clusters Add Service .
4. Select Query Processor on the **Add Service to Cluster** page and click Continue.  
The required dependencies are automatically selected.
5. Select the host on which you want to install the Query Processor by clicking View By Host. Then click Continue.

6. Select the database type, and specify the database hostname, database name, and username and password to access the database on the **Setup Database** page and click Test Connection.

**Attention:**

- (For MySQL and Oracle databases) Ensure that you specify the same name in the Database Name and Database Schemas fields.
- (For Oracle database) You must specify the value of the query\_processor\_database\_schema property (Database Schema) in upper case.

After the connection is verified, click Continue.

7. On the **Review Changes** page, accept the default settings and click Continue.

If Kerberos or Auto-TLS are set up on your Data Hub cluster, then the corresponding settings are automatically configured.

The **Command Details** page is displayed, which shows the status of the installation and configuration.

8. Click Continue if the operation is successful.
9. Click Finish on the **Summary** page.

**Results**

You are redirected to the Cloudera Manager home page. You can now see the Query Processor service listed within your cluster.

**What to do next**

(Optional) To view the **Queries** tab on the **Job Browser** page in Hue, you must enable the Query Processor Service option in Hue configurations using Cloudera Manager.



## Removing Query Processor service from cluster

You can remove the Query Processor service using Cloudera Manager when you no longer need it.

**Before you begin**

Because Hue has a dependency on the Query Processor service, deselect the Query Processor Service option in the Hue configurations before proceeding to remove service dependencies.

**Procedure**

1. Log in to the CDP Management Console as an Administrator.
2. Go to your environment and click on CM UI to open Cloudera Manager.
3. Go to Clusters Query Processor service  Delete .  
If you have not removed the service dependencies before deleting the Query Processor, then a **Remove Service Dependencies** dialog box is displayed.
  - a) Click Configure Service Dependency.
  - b) Deselect the Query Processor Service option on the **Configuration** page and click Save Changes.
  - c) Go to Clusters Query Processor service  Delete .

## Enabling the Query Processor service in Hue

The Queries tab on Hue's Job Browser page allows you to view Hive query history and query details. You need to enable the Query Processor service from the Hue configurations to enable the Queries tab.

### About this task

The Query Processor service is automatically enabled in Hue if you select it as a dependency on the **Select Dependencies** page while adding the Hue service to your cluster. But if you did not select it as a dependency while adding the Hue service, you can add the Query Processor service to your cluster (if not installed already) and then enable the Query Processor Service option in Hue configurations.

### Procedure

1. Log in to the CDP Management Console as an Administrator.
2. Go to your environment and click on CM UI to open Cloudera Manager.
3. Go to **Clusters Hue service Configuration** and select the Query Processor Service option.
4. Click **Save Changes**.
5. Restart the Hue service.

### Results

You can now view Hive query history and query details under the **Queries** tab on Hue's **Job Browser** page.

## Adding Query Processor Administrator users and groups on Data Hub clusters

Query Processor Administrators have special privileges that enable them to view and monitor queries from all users, including the ones that are submitted from query interfaces, such as Beeline, Hive Warehouse Connector (HWC), Tableau, and so on.

### About this task

This task is applicable only if you are running Hive queries.

### Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to **Clusters Query Processor Configuration** and enter the list of users in the Extra Auth Options field as follows:  
`"adminUsers": "hive, [***USER-1***], [***USER-2***]"`  
 You can also add a list of admin groups as follows:  
`"adminGroups": "admin-group, [***GROUP-1***], [***GROUP-2***]"`
3. Click **Save Changes**.
4. Restart the Query Processor service.

## Ways to clean up old queries from the Query Processor tables

Learn how to schedule a query clean-up and how to use the API to manually clean up queries from the following Query Processor tables: `vertex_info`, `dag_details`, `dag_info`, `query_details`, `hive_query`, `tez_app_info`.

### Scheduling query clean-up

By default, the Hue Query Processor automatically clears queries from the Query Processor tables that are 30 days old at 2 AM every day as per the timezone set on your server. You can configure the following parameters in Cloudera

Manager Query Processor service Configurations Query Processor Extra Configurations to set up a schedule based on your requirements:

- `hue.query-processor.event-pipeline.cleanup.cron.expression`
- `hue.query-processor.event-pipeline.cleanup-interval-secs`

For example:

```
"hue.query-processor.event-pipeline.cleanup.cron.expression" : "0 0 2 * * ?"
,
"hue.query-processor.event-pipeline.cleanup-interval-secs" : "2592000"
```

### Manually cleaning up queries using an API

The ability to clean up queries manually in addition to the scheduled clean-up routines is useful when you have a high load of queries in a particular week that are hogging resources that you must free up. The API also runs a `VACUUM` command on the Query Processor table to reclaim storage that is occupied by dead tuples.

You can send an API request using tools such as `cURL` or `Postman`.

API format: `[***QUERY-PROCESSOR-ADDRESS**]/api/admin/cleanup/[***EPOCH-TIME***]`

Where,

- `[***QUERY-PROCESSOR-ADDRESS***]` is the query processor host address
- `[***EPOCH-TIME***]` is the Unix epoch time in seconds

Queries that were run before the specified epoch time are purged.

For example:

```
curl "http://machine1.company.com:30700/api/admin/cleanup/1670006742"
```

## Downloading debug bundles

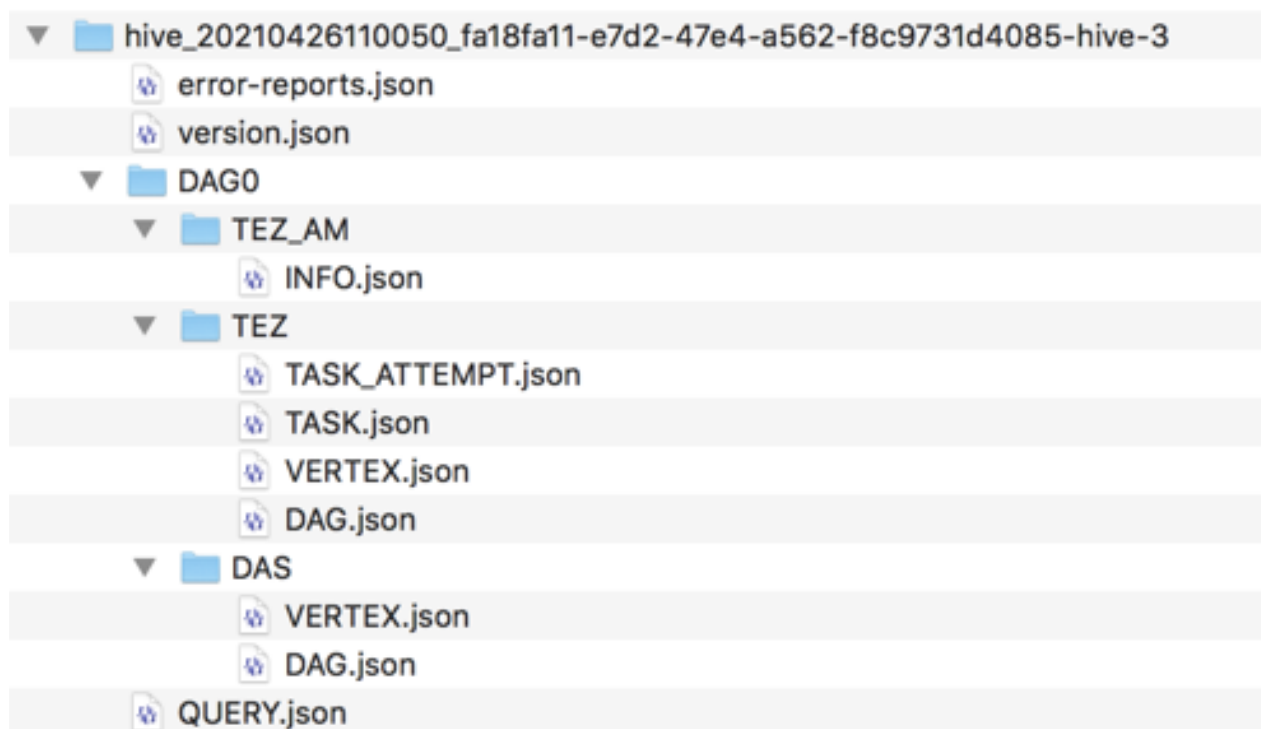
The debug bundle is a ZIP file that contains the query details in JSON format and an `error-reports.json` file, which is created only if an error occurs while the query is run.

### About this task



**Note:** This feature is available only for Hive queries.

If Tez is used to run a query, then the debug bundle also contains DAG and Tez JSON files, as shown in the following image:



### Procedure

1. Log in to the Hue web interface.
2. Click Jobs from the left assist panel.  
The **Job Browser** page is displayed.
3. Click Queries.  
The Hive queries that were run are displayed.
4. Select a query for which you want to download the debug bundle.
5. Click Download and save the ZIP file on your computer.

The filename is in the following format:

```
hive_[***HIVE-QUERY-ID***]_[***USER-ID***]_[***UNIQUE-INDEX***]
```

## Configuring Hue to handle HS2 failover

Hue can handle HiveServer2 (HS2) failovers using ZooKeeper without the need to set up a load balancer.

### About this task



**Note:** Zookeeper configuration only supports failover situations. It cannot be used to balance loads between the active HS2 servers.

### Procedure

1. Log in to Cloudera Manager as an Administrator.

2. Go to **Clusters Hue Configuration** and add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[beeswax]
hive_discovery_hs2=true
hive_discovery_hiveserver2_znode=/hiveserver2
```



**Note:** `hiveserver2` is the default value of the `hive_discovery_hiveserver2_znode` property. If this value is different from the Hive-On-Tez configuration, then you must obtain the information from the `hive.server2.zookeeper.namespace` property from Hive-On-Tez configuration.

3. Click **Save Changes**.
4. Restart the Hue service.

## Enabling Spark 3 engine in Hue

Hue leverages Apache Livy 3 to support Spark SQL queries in Hue on the Apache Spark 3 engine. To enable the Spark 3 engine, specify the Livy server URL in the Hue Advanced Configuration Snippet using Cloudera Manager, and enable the Spark SQL notebook. Livy for Spark 3 and Spark 3 services are installed when you create a Data Hub cluster with the Data Engineering cluster template.

### Before you begin



**Note:** Livy v0.6 supports Python versions upto Python 3.7. If you install Python 3.8 or higher, then you may see the following error: `TypeError: required field "type_ignores" missing from Module`.

### Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to **Clusters HDFS Configuration** and add the following lines in the Cluster-wide Advanced Configuration Snippet (Safety Valve) for `core-site.xml` field:

```
<property>
 <name>hadoop.proxyuser.hue.groups</name>
 <value>*</value>
</property>
<property>
 <name>hadoop.proxyuser.hue.hosts</name>
 <value>*</value>
</property>
<property>
 <name>hadoop.proxyuser.spark.groups</name>
 <value>*</value>
</property>
<property>
 <name>hadoop.proxyuser.spark.hosts</name>
 <value>*</value>
</property>
<property>
 <name>hadoop.proxyuser.livy.groups</name>
 <value>*</value>
</property>
<property>
 <name>hadoop.proxyuser.livy.hosts</name>
 <value>*</value>
</property>
```

3. Click **Save Changes**.

4. Go to Clusters Livy for Spark 3 service Configuration and add the following configurations:
  - a) Add the hue user in the Admin Users (livy.superusers) field.
  - b) Go to the HMS Service field and select Hive.
  - c) Click Save Changes.
5. Go to Clusters SPARK\_ON\_YARN Configuration Admin Users , add hue to the list of admin users (spark.history.ui.admin.acls) and click Save Changes.
6. Go to Clusters SPARK Configuration Admin Users , add hue to the list of admin users (spark.history.ui.admin.acls) and click Save Changes.
7. Go to Clusters SPARK 3 Configuration Admin Users , add hue to the list of admin users (spark.history.ui.admin.acls) and click Save Changes.
8. Go to Clusters Hue Configuration and enter the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini field and click Save Changes:

```
[desktop]
app_blacklist=zookeeper, pig #custom list of blocked apps
[spark]
#This is not a thrift server port
#If this TLS/SSL is enabled then check to see whether the livy url is on h
ttps or http and modify the url accordingly.
livy_server_url=http(s)://[***LIVY-FOR-SPARK3-SERVER-HOST***]:[***LIVY-
FOR-SPARK3-SERVER-PORT***]
ssl_cert_ca_verify=false
security_enabled=true
[notebook]
[[interpreters]]
[[sparksql]]
name=Spark SQL
interface=livy
```



#### Attention:

- Ensure that the Spark application is not on the blocked list.
- Set ssl\_cert\_ca\_verify=false if an SSL certificate is not present in Hue's truststore. ssl\_cert\_ca\_verify=true if your environment is a secured environment.
- Set security\_enabled=true for Kerberized clusters.

9. Restart the affected services.

### Results

You can now select the “Spark SQL” dialect on the Hue editor and run Spark queries from Hue.



**Note:** Starting a Livy session can take 30-45 seconds. If you cancel a running Spark SQL query and rerun it, the Livy session expires, and a new session is created.

## Using Hue scripts

Hue scripts help you to share workflows, set up default editors, change ownership of documents, and test the backend processes using cURL commands. The Hue scripts are included with CDP.

### Sharing workflows with users and groups

Use the following command to share workflows with users and groups:

```
./build/env/bin/hue share_all_workflows --owner [***OWNER-NAME***] --sharegr
oups [***GROUP1/USER1***],[***GROUP2/USER2***] --permissions read,write
```

To share users, use the `--shareusers` option instead of the `--sharegroups` option.

### Changing the document owner

Use the following command to change the owner of the Hue documents:

```
./build/env/bin/hue change_owner_of_docs --olduser [***OLD-USER-NAME***] --newuser [***NEW-USER-NAME***]
```

### Testing the backend using cURL

Use the following command to get the syntax for the backend cURL commands or to test the endpoints:

```
./build/env/bin/hue get_backend_curl --showcurl --testname [***SERVICE-NAME-WHERE-TEST-NEEDS-TO-BE-PERFORMED***] --service [***COMMA-SEPARATED-SERVICES-TO-TEST***]
```

Replace `[***SERVICE-NAME***]` with the name of the service where you want to perform the test. Replace `[***COMMA-SEPARATED-SERVICES-TO-TEST***]` with options or services such as `all`, `https`, `solr`, `oozie`, `rm`, `jhs`, `sparkhs`, and so on.

### Setting the default editor in Hue

Use the following command to set the default editor in Hue:

```
./build/env/bin/hue set_default_editor --hive/impala --username [***USERNAME***]
```

The default value for the username is `all`.

## Configurations for submitting a Hive query to a dedicated queue

When you run a Hive query from Hue or Beeline, HiveServer2 submits the query in YARN in the queue mapped from an end user instead of a hive user. If you run a query while specifying `doAs=false`, then the queries are submitted and run as a hive user, using impersonation. This is the default behaviour. You can change this behaviour by configuring the queue using `"yarn_queue_manager"` and `"application tag"` parameters in YARN.

If you see go to the **Job Browser** page in Hue and see that the Hive queries are appearing under the default queue, and not the actual user who ran the query, then this means that the queries are not being run in their dedicated queue in YARN. To submit a Hive query in its dedicated queue, you must enable application-tag based scheduling by configuring queue mapping and per queue properties.

### Related Information

[Configuring queue mapping to use the user name from the application tag using Cloudera Manager](#)

[Configure Per Queue Properties](#)

## Configuring timezone for Hue

You can configure the timezone for Hue using Cloudera Manager. The default value is `America/Los_Angeles`.

### Procedure

1. Log in to Clouder Manager as an Administrator.

2. Go to `Clusters Hue Configuration Time Zone` and set a desired value.  
The default value is “America/Los\_Angeles”. You can refer to <https://timezonedb.com/time-zones> for information on the timezones.
3. Click `Save Changes`.
4. Restart the Hue service.