Cloudera Runtime 7.2.16

Administering Hue

Date published: 2020-08-10
Date modified: 2022-12-13

CLOUDERA

Legal Notice

Contents

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------- 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 Cloudera Data Warehouse:

Parameter	Description
``` [[desktop]] app_blacklist ```	Used to add or remove applications, such as the File Browser, Impala, Hive, Oozie, and so on from the blocked list of applications.

``
[desktop]
    [[session]]
    ttl=[***NUMBER-OF-SECONDS***]
``
	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.

``` [jobbrowser] [[query_store]] server_url=[***QUERY-PROCESSOR-URL***] ```	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 : ``` [jobbrowser] [[query_store]] is_enabled=false ``` To enable the **Queries** tab, set is_enabled to true.
``` [aws] [[aws_accounts]] [[[default]]] access_key_id=[***AWS-ACCESS-KEY***] secret_access_key=[***SECRET-ACCESS-KEY***] region=[***AWS-REGION***] ```	Used to enable the S3 File Browser for Hue without IDBroker.
``` [aws] has_iam_detection=true [[aws_accounts]] [[[default]]] region=[***AWS-REGION***] ```	Used to enable the S3 File Browser for Hue with IDBroker.

``
[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/
``
	Used to enable the ABFS File Browser for Hue without IDBroker.

``` [desktop] [azure] ```	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.

Environment safety valves

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

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
```


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: • 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
runcpserver.log	Hue (CherryPy) web server info (CP server runs Django core)
hue_install.log	Contains the log produced during installation

Enabling DEBUG

DEBUG is available for the Hue Django 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.

About this task

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

Note: Django logs every SQL query it executes in the debug mode. This could rapidly increase memory consumption.

You can enable DEBUG messages for all the logs in var/log/hue from Cloudera Manager or from the Hue web UI.

Using Cloudera Manager

1. Go to Cloudera Manager Clusters Hue Configuration .
2. Select the Enable Django Debug Mode option.
3. Click Save Changes.
4. Restart the Hue service.

Using the Hue web interface

1. Log into the Hue web interface as a superuser.
2. Click Administer Server from the left-assist pannel.

   About Hue page is displayed.
3. Click on the Server Logs tab.
4. Select the Force DEBUG level option.

   Debug is enabled on-the-fly.

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

For example:

```
CustomLog "|/usr/sbin/rotatelogs /var/log/hue-httpd/access_log.%Y-%m-%d-
%H_%M_%S 5B" common
ErrorLog "|/usr/sbin/rotatelogs /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/rotatelogs -l /var/log/hue-httpd/access_log.%Y-%m-
%d-%H_%M_%S 86400" common
ErrorLog "|/usr/sbin/rotatelogs -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 *rotatelogs - Piped logging program to rotate Apache logs*.

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

Related Information

rotatelogs - 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.

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 test 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-4.05
017,-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 18.41163,23.95773c1.518
1,-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 18.4116
4,23.95773c1.5181,-1.36927
22.76959,-12.43803 18.41289,-23.96533z" fill="#FF5A79"/> <path stroke="nu
ll" 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:
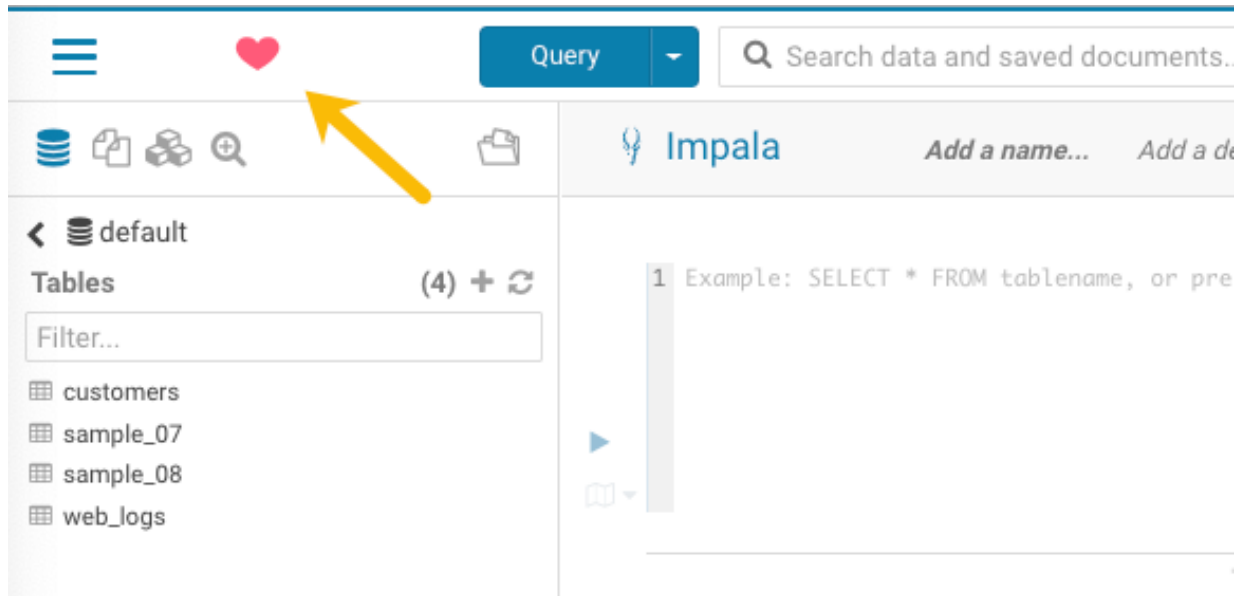
8.

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.
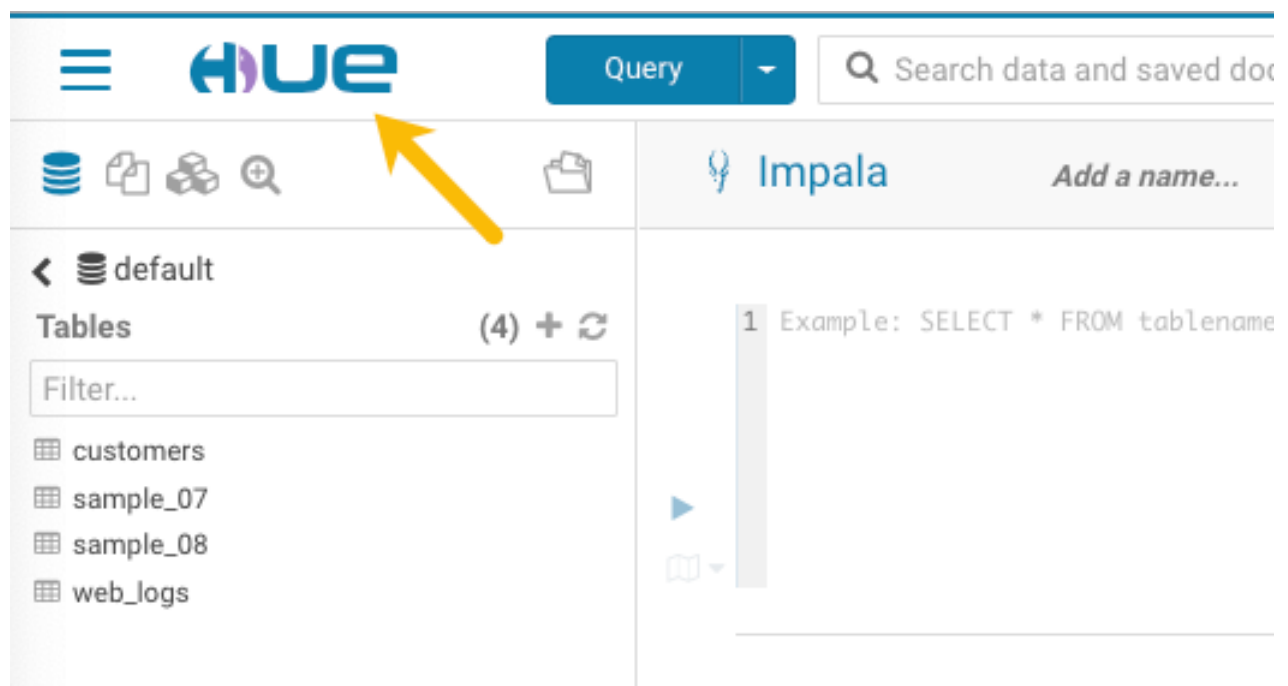
Procedure

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 date 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 -m1 python|awk '{ print
    $2 }')/environ|egrep -v "^HOME=|^TERM=|^PWD="`;do export $line;done
   ```

   • Ubuntu:

   ```
   for line in `strings /proc/$(lsof -i :8888|grep -m1 hue|awk '{ print $2
    }')/environ|egrep -v "^HOME=|^TERM=|^PWD="`;do export $line;done
   ```

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_PAS SWORD 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 passw
ord> /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/bu
ild/env/bin/hue changepassword admin
```

- Promote Hue user to superuser

```
HUE_IGNORE_PASSWORD_SCRIPT_ERRORS=1 /opt/cloudera/parcels/CDH/lib/hue/bu
ild/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/bu
ild/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

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

⚠ **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.

Before you begin

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

This task assumes that you already have a PostgreSQL database installed on a host in your cluster.

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 psql terminal by entering the following command:

```
sudo -u postgres psql
```

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

```
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***];
```

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.

   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.

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.

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.

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 Environment Advanced Configuration Snippet (Safety Valve) 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 Spaqrl SQL notebook. Livy for Spar 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>
<name>hadoop.proxyuser.hue.hosts</name>
 <value>*</value>
 <name>hadoop.proxyuser.spark.groups</name>
 <value>*</value>
 <name>hadoop.proxyuser.spark.hosts</name>
 <value>*</value>
 <name>hadoop.proxyuser.livy.groups</name>
 <value>*</value>
 <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
#To enable the Notebook editor on the left-assist panel.
[notebook]
show_notebooks=true
[[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.
> - Enabling the Notebook editor on the left-assist panel is optional.

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***] --n
ewuser [***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, httpfs, 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 as described in Configuring queue mapping to use the user name from the application tag using Cloudera Manager and 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.