

Cloudera Runtime 7.1.6

Hue Troubleshooting

Date published: 2020-07-28

Date modified: 2021-03-19

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

The Hue load balancer not distributing users evenly across various Hue servers.....	6
Unable to authenticate users in Hue using SAML.....	6
Cleaning up old data to improve performance.....	7
Unable to connect to database with provided credential.....	8
Activating Hive query editor on Hue UI.....	9
Query execution finished in Hue but shows as executing on Cloudera Manager Impala Queries Page.....	9
Finding the list of Hue superusers.....	10
Username or password is incorrect while accessing Hue through Knox.....	11
HTTP 403 error while accessing Hue UI from Knox.....	12
Unable to access Hue from Knox Gateway UI.....	14
Referer checking failed because domain does not match any trusted origins.....	16
Unable to view Snappy-compressed files.....	17
“Unknown Attribute Name” exception when enabling SAML.....	18
Impala query fails with invalid query handle error.....	19
Services backed by PostgreSQL fail or stop responding.....	20

Error validating LDAP user in Hue.....	21
502 Proxy Error while accessing Hue from the Load Balancer.....	22
Invalid method name: 'GetLog' error after submitting Hive queries.....	22
"Authorization Exception" error on submitting queries in Hue.....	23
Cannot alter compressed tables in Hue.....	25
Connection failed error when accessing the Search app (Solr) from Hue.....	25
Downloading query results from Hue takes time.....	26
Hue Load Balancer does not start after enabling TLS.....	26
Unable to terminate Hive queries from the Hue Job Browser in a Kerberized cluster.....	27
Unable to view or create Oozie workflows in Hue on a Knox-secured cluster.....	28
MySQL: 1040, 'Too many connections' exception.....	28
Unable to connect Oracle database to Hue using SCAN.....	29
Increasing the maximum number of processes for Oracle database.....	30
UnicodeDecodeError: UTF-8 codec cannot decode byte in position.....	31
UnicodeDecodeError: 'ascii' codec cannot decode byte in position.....	31
Fixing authentication issues between HBase and Hue.....	32
Hue Load Balancer does not start due to lengthy BalancerMember Route length.....	33

Enabling access to HBase browser from Hue.....	34
Fixing a warning related to accessing non-optimized Hue.....	35
Hue load balancer does not start after enabling TLS.....	35

The Hue load balancer not distributing users evenly across various Hue servers

The Hue load balancer redirects the new users to the newly added Hue servers and the existing users to the existing Hue servers on your cluster. Even though you add more Hue servers to meet the growing user base, the resources might not be utilized effectively.

About this task

The Hue load balancer is tasked to evenly distribute users across the available Hue servers for effective resource utilization. However, due to session persistence, it does not distribute users evenly. To overcome this issue, you can refresh the cookies from Cloudera Manager.

The load balancer uses the cookie ROUTEID with a random string from your browser and is stored in the hue.conf file. This random string is used to redirect a user to the Hue server. To refresh the cookie and set a new random string every time you add a new Hue server, do the following:

Procedure

1. Go to Cloudera Manager Clusters Hue Configuration .
2. Click Scope Load balancer and select the Hue Load Balancer Cookie Refresh checkbox.
This refreshes the cookie value in the hue.conf file to rebalance the Hue backend connections.
3. In the Instances tab, select all the Hue services and roles, and click Action for Services Restart .
This creates a new random string for the cookie which the load balancer can now use to evenly distribute users.

Results

Whenever you restart the Hue server, the load balancer will redistribute users evenly based on the server load.

Unable to authenticate users in Hue using SAML

If you have configured SAML to authenticate users, but your users are unable to log into Hue using Single Sign On (SSO), then it is possible that the RSA key format is not supported. To resolve this issue, you can use an unprotected private key and then specify the private key filename in the safety valve.

Procedure

1. Convert the .key file to an unprotected private key file by using the following command:

```
openssl rsa -in /opt/cloudera/security/<file name>.key -out /opt/cloudera/security/<file name_unprotected>.key
```

```
openssl rsa -in /opt/cloudera/security/hadoop-cpi-prod.key -out /opt/cloudera/security/hadoop-cpi-prod_unprotected.key
```

2. Update the advanced configuration snippet as shown in the following example:

```
[libsaml]  
xmlsec_binary=/usr/bin/xmlsec1  
metadata_file=/opt/cloudera/security/saml/idp-openam-metadata.xml  
key_file=/opt/cloudera/security/hadoop-cpi-prod_unprotected.key  
cert_file=/opt/cloudera/security/hadoop-cpi-prod.pem
```

Cleaning up old data to improve performance

Some tables in Hue retain data indefinitely resulting in slower performance or application crash. Hue does not automatically clean up data from these tables. You can configure Hue to retain the data for a specific number of days and then schedule a cron job to clean up these tables at regular intervals for improved performance.

About this task

Consider cleaning up old data from the backend Hue database if you face the following problems while using Hue:

- Upgrade times out
- Performance is slower than expected
- Long time to log in to Hue
- SQL query shows a large number of documents in tables
- Hue crashes while trying to access saved documents

Before you begin

Back up your database before starting the cleanup activity. Check the saved documents such as Queries and Workflows for a few users to prevent data loss. You can also note the sizes of the tables you want to clean up as a reference by running the following queries:

```
select count(*) from desktop_document;  
select count(*) from desktop_document2;  
select count(*) from beeswax_session;  
select count(*) from beeswax_savedquery;  
select count(*) from beeswax_queryhistory;  
select count(*) from oozie_job;
```



Note: The optimal number of documents that can be stored in a table is less than or equal to 30,000. Consider this number while specifying the cleanup interval.

Procedure

1. SSH in to an active Hue instance.
2. Change to the Hue home directory:

```
cd /opt/cloudera/parcels/CDH/lib/hue
```

3. Run the following command as the root user:

```
DESKTOP_DEBUG=True ./build/env/bin/hue desktop_document_cleanup --keep-days x --cm-managed
```

The `--keep-days` property is used to specify the number of days for which Hue will retain the data in the backend database.

```
DESKTOP_DEBUG=True ./build/env/bin/hue desktop_document_cleanup --keep-days 30 --cm-managed
```

In this case, Hue will retain data for 30 days.

The logs are displayed on the console because `DESKTOP_DEBUG` is set to `True`. Alternatively, you can view the logs from the following location:

`/var/log/hue/desktop_document_cleanup.log`

The first run can typically take around 1 minute per 1000 entries in each table.

4. Check whether the table size has decreased by running a query as follows:

```
select count(*) from desktop_document;
```

If the `desktop_document_cleanup` command has run successfully, the table size should decrease.

What to do next

Set up a cron job that runs at regular intervals to automate the database cleanup. For example, you can set up a cron job to run daily and it purges data older than *x* number of days.

Unable to connect to database with provided credential

Cloudera Manager tests the database connection when you add the Hue service to a cluster. The “Test Database Connection” does not work for Oracle databases that require service name instead of the Oracle System ID (SID). This could stop you from adding the Hue service to your cluster.

About this task

If you encounter the following error while adding the Hue service through Cloudera Manager, then follow the workaround as mentioned in this topic:

Unable to connect to database with provided credential. Able to find the Database server, but not the specified database. Please check if the database name is correct and make sure that the user can access the database.

Procedure

1. Install a different database instance, such as MySQL temporarily to use with Hue.
This is referred to as the Hue database.
2. Add the Hue service from Cloudera Manager and specify the Hue database details that you created in the previous step.

This will allow you to get past the Add Service wizard and add the Hue service to your cluster.

3. Modify the Hue instance to use the actual Oracle database as follows:
 - a) Navigate to Cloudera Manager Clusters Hue service Configuration Category Database .
The database configuration fields are displayed.
 - b) Set up the Oracle database by configuring the following fields:
 1. Select Oracle as the Hue Database Type.
 2. In the Hue Database Hostname field, specify the Fully Qualified Domain Name (FQDN) of the host on which you have installed the Oracle database.
 3. In the Hue Database Port field, specify the port on the host on which the Oracle database is running.
Typically, this value is 1521.
 4. In the Hue Database Username field, specify the username to log in to the Oracle database.
 5. In the Hue Database Password field, specify the database password.
 6. In the Hue Database Name field, specify the name of the Hue database in the following format:

```
<HUE_DB_HOST>:1521/<servicename>
```

- c) Navigate to Cloudera Manager Clusters Hue service Configuration Category Advanced and specify the following in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[desktop]
[[database]]
port=0
```

4. Click Save Changes.

5. Restart the Hue service by clicking **Actions Restart**.

Activating Hive query editor on Hue UI

You may not see the Hive query editor on the Hue user interface if you have not installed and selected the HIVE_ON_TEZ service on your cluster. The HIVE_ON_TEZ service is needed to configure and use Hive with Hue.



Note:

In CDH 6 and earlier, the Hive service included the Hive Metastore and HiveServer2. In Cloudera Runtime 7.0 and later, this service includes only Hive Metastore. HiveServer2 and other components of the Hive execution engines are part of the HIVE_ON_TEZ service.

About this task

To enable the Hive query editor on the Hue web UI:

Procedure

1. Sign in to Cloudera Manager as an Administrator.
2. Check whether the HIVE_ON_TEZ service is installed on your cluster.
If it is not installed already, add it as a service.
3. Go to **Clusters Hue service Configuration**.
The list of Hue configurations is displayed.
4. Search for the HiveServer2 Service field and select the HIVE_ON_TEZ service.
5. Save the changes and restart Hue.

Results

The Hive query editor should now be available on the Hue user interface.

Query execution finished in Hue but shows as executing on Cloudera Manager Impala Queries Page

Cloudera Manager and the Impala demon web page may show a query in an “executing” or “In Flight” state even though the query has finished executing on the Hue web UI. This can happen due to various reasons.

The three main reasons why the completed Hue query still shows as “executing” are:

- Hue does not close the connection to Impala until you click on the **Results** page.
Clicking the **Results** page in Hue executes the fetchresults call to Impala.
- Impala queries are client-driven. Therefore, the query still remains in a running state until the client sends a fetch command to complete fetching the entire result set.
- If a query has not been closed or unregistered, Impala shows the same in the **In Flight** section on its web UI. Cloudera Manager shows all In Flight queries in the “Executing” state.

Impala query life cycle

When you submit Impala queries, they are first registered by the system. The system identifies the queries with the help of a coordinator. They also have a state, such as CREATED, INITIALIZED, RUNNING, FINISHED, EXCEPTION, and some metadata.

- FINISHED implies that the rows are available but not all rows are ready to be fetched. It is possible that Impala daemons are still executing the query.
- EXCEPTION implies that an error has occurred. For example, if the system runs out of memory, then the query transitions to the EXCEPTION state.

The query can also go into an EXCEPTION state if it is cancelled.

Query cancellations may be triggered explicitly with a HiveServer2/Beeswax call or if the query times out. Query time-out may be set through a process-wide `impalad` argument or with a per-query option.

Currently, Impala does not have a state that explicitly indicates whether all Impala daemons have finished executing the query and that all results have been fetched. Let us call it as End of Statement (EOS), temporarily.

When a query is in the EOS (FINISHED) or EXCEPTION state, the query is not doing any more processing, but the query remains registered. It needs to remain registered because clients may need to access the state.

The query is unregistered only in the following two cases:

- The query is explicitly closed by a `Close()` API call
- The session associated with the query is closed explicitly or the session time-out is set and the session times out



Note: Hue does not close a query until you explicitly close it. When you close a browser tab on which you are running the query in Hue, the browser sends a JavaScript `Close()` callback request to close the query. If you leave the query unattended, for example by closing the laptop through which you are accessing Hue or if the browser crashes, then the `Close()` call is never sent to Hue. The query may eventually time out, but because it was not cancelled explicitly, it does not properly clean up the resources.

To optimize resource utilization, configure the Impala daemon to kill the idle sessions by setting the session timeout value in the `--idle_session_timeout` `impalad` argument:

1. Sign in to Cloudera Manager as an Administrator.
2. Go to Clusters Impala service Configuration .
3. Specify the following in the Impala Command Line Argument Advanced Configuration Snippet (Safety Valve) field:

```
--idle_session_timeout=<maximum lifetime of your queries in seconds>
```

For example,

```
--idle_session_timeout=3600
```

In this case, the query will time out after one hour.

Finding the list of Hue superusers

You can fetch the list of superusers by using the Hue shell with Python code or by running a SQL query on the `auth_user` table.

Using the Hue shell and Python code to find Hue superusers

1. Connecting to Hue shell by running the following command:

```
/opt/cloudera/parcels/CDH/lib/hue/build/env/bin/hue shell --cm-managed
```

2. Enter the Python code as follows:

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

```
print "%s" % User.objects.filter(is_superuser = True)
```

Sample output:

```
<QuerySet [<User: admin>]>
```

Running a SQL query on the auth_user table to find Hue superusers

1. Connect to Hue database shell by running the following command:

```
/opt/cloudera/parcels/CDH/lib/hue/build/env/bin/hue dbshell --cm-managed
```

2. Run the following SQL query:

```
select username, is_superuser from auth_user where is_superuser=1;
```

Sample output:

```
-----+
username is_superuser
-----+

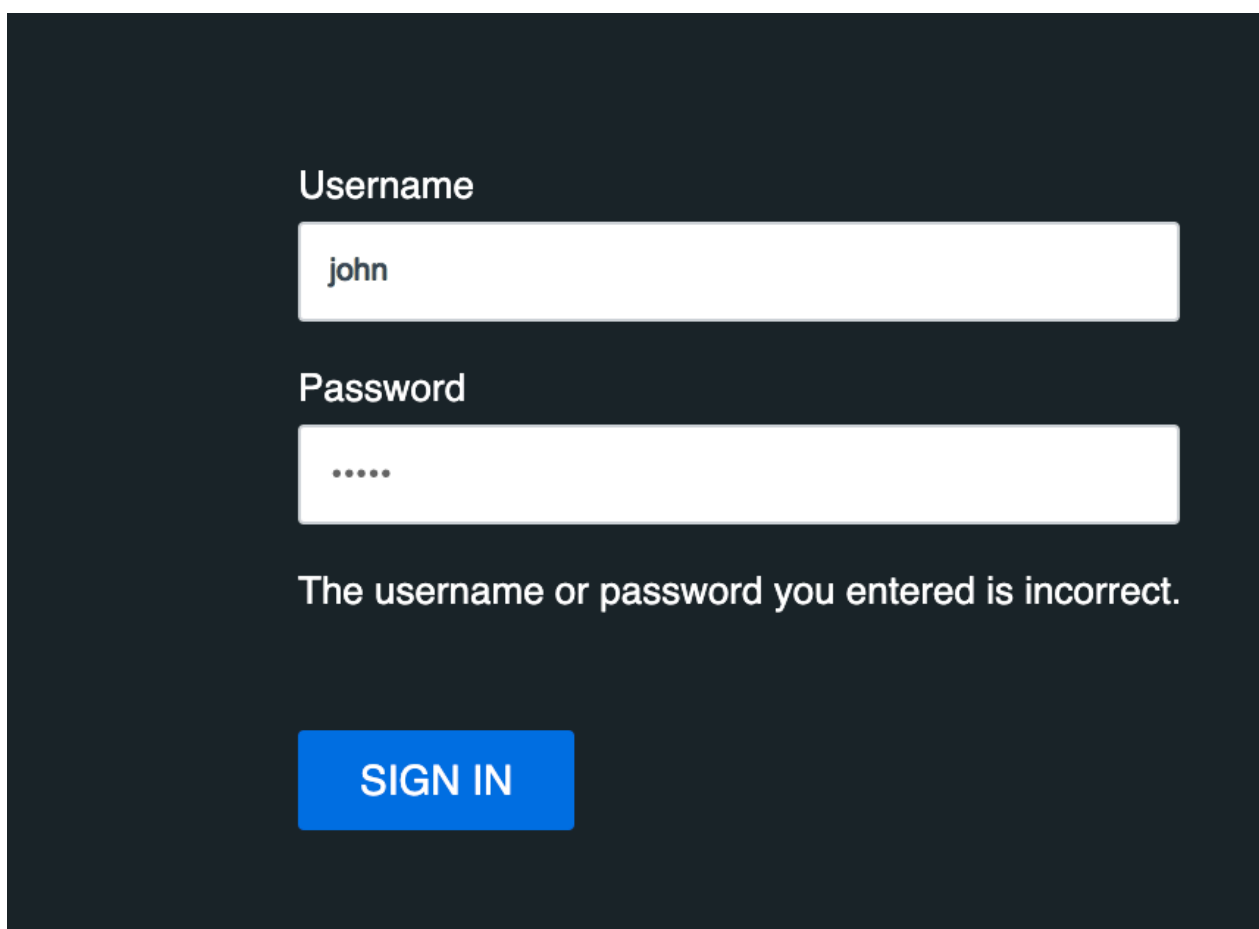
admin 1
-----+
1 row in set (0.00 sec)
```

Username or password is incorrect while accessing Hue through Knox

If an error such as “The username or password you entered is incorrect” appears when you try to log in to Hue using the Knox UI, then you can verify your credentials by logging in to the Knox Gateway using the command-line interface.

About this task

Figure 1: Knox Gateway UI: Incorrect username or password



The screenshot shows a dark-themed login interface. At the top, the label "Username" is followed by a text input field containing the text "john". Below this, the label "Password" is followed by a password input field with five dots. A message in the center states "The username or password you entered is incorrect." At the bottom, there is a blue button with the text "SIGN IN" in white capital letters.

Procedure

1. Open a terminal session.
2. SSH into the Knox Gateway host by entering the following command:

```
ssh [***KNOX-USERNAME***]@[***KNOX-HOST***].[***DOMAIN***].site
```

Replace *KNOX-HOST.domain.site* with the Fully-Qualified Domain Name (FQDN) of your Knox Gateway host.
For example:

```
ssh john@abc-1.example.com
```

3. Enter the password that you used on the Knox Gateway web UI.
If you are able to log in using these credentials, then you should be able to log into the Knox Gateway UI.

HTTP 403 error while accessing Hue UI from Knox

When an HTTP 403 error appears while accessing the Hue UI from the Knox Gateway, then the user or the group that the user belongs to may not have the required permissions.

Procedure

1. Sign in to Cloudera Manager as an Administrator.

2. Go to ClustersRanger service Ranger Admin Web UI and sign in to the Ranger Admin web UI as an Administrator.
3. Click on cm_knox.
4. The cm_knox policies are displayed.

Service Manager > cm_knox Policies

List of Policies : cm_knox

Search for your policy...

Add New Policy

Policy ID	Policy Name	Policy Labels	Status	Audit Logging	Roles	Groups	Users	Action
16	all - topology, service	--	Enabled	Enabled	--	--	admin, rangerlookup	
17	CDP Proxy UI and API	--	Enabled	Enabled	--	public	--	

5. Click on the policy ID 16 (all - topology, service).
The **Edit Policy** page is displayed.

Service Manager > cm_knox Policies > Edit Policy

Edit Policy

Policy Details :

Policy Type: **Access**

Policy ID: **16**

Policy Name *: all - topology, service

Policy Label: Policy Label

Knox Topology *: * **include**

Knox Service *: * **include**

Description: Policy for all - topology, service

Audit Logging: **YES**

- Verify that the user and the group to which the user belongs has the required permissions.

If the user or group does not have the required permission, then add the user or the user group to the Select User field and select Allow under Permissions.

Allow Conditions :

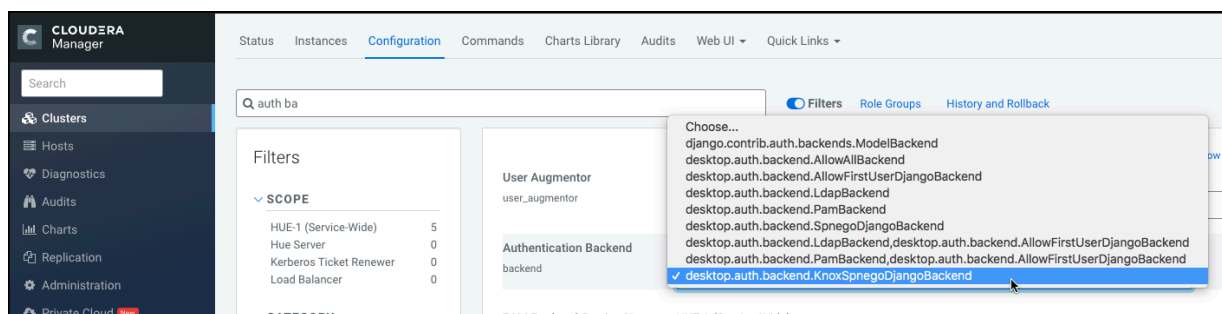
Select Role	Select Group	Select User	Policy Conditions	Permissions	Delegate Admin	
Select Roles	Select Groups	x admin	Add Conditions +	Allow	<input checked="" type="checkbox"/>	x
Select Roles	Select Groups	x rangerlookup	Add Conditions +	Allow	<input type="checkbox"/>	x
Select Roles	Select Groups	x hue	Add Conditions +	Allow	<input type="checkbox"/>	x

Unable to access Hue from Knox Gateway UI

If you are unable to access Hue from the Knox Gateway UI, then it is possible that the KnoxSpnegoDjangoBackend property is not configured for Hue. Cloudera Manager does not set the Authentication Backend to KnoxSpnegoDjangoBackend by default even though Knox is enabled on your cluster.

Procedure

- Sign in to Cloudera Manager as an Administrator.
- Go to Clusters Hue service Configurations and search for the Authentication Backend field.
- Select desktop.auth.backend.KnoxSpnegoDjangoBackend from the dropdown.




- Click Save Changes.
- Go to Clusters \$Knox service Instances and note down the hostnames of the Knox Gateways.

You must provide these details in the next step.

If you have set up Knox in High-Availability (HA) mode, then you can see more than one Knox Gateways listed on the Instances tab.

- Go back to Clusters Hue service Configurations and search for the Knox Proxy Hosts field.
- Enter the hostname of the Knox Gateway that you noted earlier.

If you have set up Knox HA, then click + to add another hostname.




8. Click Save Changes.


You would see the following warning:

Role is missing Kerberos keytab. Go to the Kerberos Credentials page and click the Generate Missing Credentials button.


9. Click Administration on the Cloudera Manager left navigation panel and select Security.
10. Go to the Kerberos Credentials tab and click Generate Missing Credentials.



A pop-up showing the status is displayed.




11. Go to Clusters Hue service and click Restart next to Actions.



12. On the **Stale Configurations** page, click Restart Stale Services.
The **Restart Stale Services** wizard is displayed.
13. On the Review Changes page, select Redeploy client configuration, and click Restart Now.
The **Command Details** page shows the live status as the service restarts.
When all the steps are complete, click Finish.

14. From the Hue service page, click **Web UI Knox Gateway UI**.




The Knox Gateway UI is displayed.

15. On the **General Proxy Information** page, expand the CDP Proxy topology by clicking + cdp-proxy under Topologies.

The list of services that are configured with the cdp-proxy topology is displayed.

16. Click on the Hue logo.



You should be able to log in to the Hue web UI.

You can also log into Hue using the following URL:

```
https://[***HOSTNAME***]:[***PORT***]/gateway/cdp-proxy/hue/
```

Referer checking failed because domain does not match any trusted origins

You may see a "Referer checking failed" error in the Hue access.log file if the Knox Gateway DNS does not resolve properly when accessing Hue using the Knox Gateway UI.

About this task

Follow the steps listed in this topic if you see the following error in the Hue access.log file:

```
"POST /accounts/login HTTP/1.1" - Referer checking failed -https://<ip_address>:<knoxui_port>/gateway/cdp-proxy/hue/hue/accounts/login?next=%2F%253FdoAs%253Dknoxui does not match any trusted origins."
```

Before you begin

The Knox Gateway DNS may not resolve due to various reasons. Verify the following before proceeding:

- Verify that the Knox proxy user configuration is set properly in the Hue configuration
- Check whether the hostnames specified in the Hue service Configuration Knox Proxy Hosts fields are spelled correctly
- Make sure that the DNS is set up correctly by logging into the DNS server and pinging a few hosts
- Make sure that the /etc/hosts file has the correct mapping of the IP address and the hosts

Procedure

1. Go to Clusters \$Knox service Instances and click on the Knox Gateway hostname.
2. Note down the IP address of the Knox Gateway host from the Details section.
3. Go to Clusters Hue service Configurations and search for the Knox Proxy Hosts field.
4. Enter the IP address of the Knox Gateway host that you noted earlier.
5. Click Save Changes.
6. Verify whether you can access the Knox Gateway through the following URL:

```
https://[***IP_ADDRESS***]:[***KNOXUI_PORT***]/gateway/cdp-proxy/hue/hue/accounts/login?
```

Unable to view Snappy-compressed files

You must install the python-snappy library on your cluster to view files compressed with Snappy using the Hue File Browser and the HBase Browser. Post-installation, Hue automatically detects and displays the Snappy-compressed files.

Before you begin

The python-snappy library is incompatible with the python library called snappy. You must uninstall snappy if it is present on your cluster.

Run the following command to check whether the snappy library is installed on your cluster:

```
/usr/bin/pip show snappy
```

No output on the console indicates that the snappy library is not installed on your cluster. If you get any results for snappy, then uninstall it by running the following command:

```
/usr/bin/pip uninstall snappy
```

Next, check whether you have the python-snappy library is installed on your cluster by running the following command:

```
/usr/bin/pip show python-snappy
```

Sample output:

```
Name: python-snappy
Version: 0.5.4
Location: /usr/lib64/python2.7/site-packages
```

Procedure

1. Sign in to Cloudera Manager as an Administrator.
2. Stop the Hue service by going to Cluster Hue service Action .

3. Change to the following directory depending on whether you have used parcels or packages to set up your CDH cluster.

For parcels:

```
cd /opt/cloudera/parcels/CDH/lib/hue
```

For package:

```
cd /usr/lib/hue
```

4. Install the python-snappy package by running the following commands:

```
yum install gcc gcc-c++ python-devel snappy-devel  
./build/env/bin/pip install -U setuptools  
./build/env/bin/pip install python-snappy
```

5. Verify that the python-snappy library is readable by all users by running the following commands:

```
ls -lart `locate snappy.py`
```

The output should be similar to the following:

```
-rw-r--r-- 1 root root 11900 Sep  1 12:25 /usr/lib64/python2.7/site-pack  
ages/snappy.py  
-rw-r--r-- 1 root root 10344 Sep  1 12:26 /usr/lib64/python2.7/site-packa  
ges/snappy.pyc
```

6. Start the Hue service by going to Cluster Hue service Action .
7. Verify that the python-snappy library is working for Hue by running the following command:

```
sudo -u hue /bin/bash -c "echo 'import snappy' | python"
```

If the python-snappy library is working as expected, then no output is displayed for this command.

Results

You should be able to view Snappy-compressed files on the Hue File Browser and the HBase Browser using the Hue web interface.

“Unknown Attribute Name” exception when enabling SAML

You may see an “Unknown Attribute Name” exception when a SAML Identity Provider (IdP) returns the 'uid' profile attribute and Hue which uses pysaml2 cannot interpret this attribute. To resolve this, you must create an attribute mapping file and then reference it in the libsaml configuration of Hue.

Procedure

1. SSH into a Hue server as a root user.
2. Create an attribute mapping directory as follows:

```
mkdir -p /opt/cloudera/security/saml/attribute_mapping
```

3. Create an attribute mapping file as follows:

```
vi /opt/cloudera/security/saml/attribute_mapping/saml_uri.py
```

4. Add the following lines in the `saml_uri.py` file:

```
MAP = {
    "identifier": "urn:oasis:names:tc:SAML:2.0:attrname-format:uri",
    "fro": {
        'uid': 'uid',
    },
    "to": {
        'uid': 'uid',
    }
}
```

5. Repeat steps 1 to 4 on all the Hue hosts.
6. Sign in to Cloudera Manager as an Administrator.
7. Go to Clusters Hue service Configuration .
8. Add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[libsaml]
xmlsec_binary=/usr/bin/xmlsec1
metadata_file=/opt/certs/saml/FederationMetadata.xml
key_file=/opt/certs/hue.key
cert_file=/opt/certs/hue.crt
entity_id=[**HOST-BASE-URL**]
logout_enabled=false
username_source=attributes
attribute_map_dir=/opt/cloudera/security/saml/attribute_mapping
#user_attribute_mapping='{ "uid": "username" }'
```

9. Click Save Changes.
10. Restart the Hue service by clicking Actions Restart .

Results

The users should now be able to authenticate to Hue through SAML.

Impala query fails with invalid query handle error

You encounter an "Invalid query handle" error when running Impala queries from the Hue web interface because the connection between Impala Thrift server and the Hue Load Balancer times out. This is governed by the `server_conn_timeout` property.

About this task

The default value of the `server_conn_timeout` property is 30 minutes. You can increase the timeout limit by updating the Hue configuration using Cloudera Manager.

Procedure

1. Log into 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. Increase the value of the `server_conn_timeout` property in the Impala section as follows:

```
[impala]
server_host=[**SERVER-HOST**]
server_port=[**PORT**]
```

```
server_conn_timeout=[ ***TIMEOUT-IN-SECONDS*** ]
```

You can increase the value of the `server_conn_timeout` property to 2 hours (7200 seconds).

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

Results

You should be able to run Impala queries from the Hue web interface successfully.

Services backed by PostgreSQL fail or stop responding

When the number of connections between the CDP services and the PostgreSQL database exceeds the preset connection limit, it is possible that a new connection fails, Cloudera Manager stops responding, and you cannot log into Hue. The logs show a “FATAL: remaining connection slots are reserved for non-replication superuser connections” error.

About this task

The number of connections between the CDP services and the PostgreSQL database is governed by the `max_connections` setting. By default, the maximum number of available connections to your PostgreSQL database is 115. 15 connections are reserved for the superuser to maintain the state and integrity of your database, and 100 connections are available for CDP and other services.



Note: There is a higher probability for this issue to occur in Hue because Cloudera Manager starts the Hue service after starting other CDP services. Therefore, the Hue service gets relatively fewer connections to PostgreSQL as compared to other services sharing the same database.

Procedure

1. Check the number of available and idle connections:
 - a) SSH into the PostgreSQL database from the command-line client `psql` as an admin user.
 - b) Run the following query to check the number of idle connections:

```
SELECT datname, count(datname) FROM pg_stat_activity WHERE state = 'idle' GROUP BY datname;
```

- c) Run the following query to check the number of connections currently in use:

```
SELECT datname, count(datname) FROM pg_stat_activity GROUP BY datname;
```

- d) Run the following command to view the maximum number of connections:

```
show max_connections;
```

- e) Run the following query to know where the connections are going:

```
SELECT datname, numbackends FROM pg_stat_database;
```

2. If most connections are idle and the `max_connections` value is less than 100, then increase the `max_connections` value in the `postgresql.conf` file:
 - a) Log into Cloudera Manager and stop all services that use the PostgreSQL database.
 - b) SSH into the host on which the PostgreSQL server is running.
 - c) Open the `postgresql.conf` file for editing.

The `postgresql.conf` file is typically present in the `/var/lib/pgsql/data` directory. But this may vary depending on where you have installed the database.
 - d) Increase the value of `max_connections` as per the following recommendation:

Allow a maximum of 100 connections for each database and add 50 extra connections. For example, for two databases, set the maximum connections to 250.

If you store five databases on one host (the databases for Cloudera Manager Server, Activity Monitor, Reports Manager, Apache Atlas, and Hive Metastore), set the maximum connections to 550.
 - e) Save the changes and exit.
 - f) Restart the PostgreSQL database by running the following command:
- ```
pg_ctl restart
```
- g) Restart all the affected services from Cloudera Manager.

### What to do next

If increasing the connection limit does not solve your problem and you see a need to scale up, then add new PostgreSQL instances on other hosts and migrate the services to those hosts with the help of your Database Administrator (DBA).

## Error validating LDAP user in Hue

Hive can use LDAP and Kerberos both, if they are enabled on your CDP cluster. By default, Hive uses LDAP to authenticate the Hue service instead of Kerberos. As a result, you may see the following error after logging in to Hue web interface or while trying to access the Hive editor: Bad status: 3 (PLAIN auth failed: Error validating LDAP user).

### About this task

Additionally, you may not be able to view databases or Hive tables. To resolve this issue, you can enforce the client connections (between Hive and Hue) to use Kerberos instead of LDAP by configuring the value of the `hive.server2.authentication` property to `KERBEROS` in the Hue `hive-site.xml` file.

### Procedure

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration Hue Server Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`.
3. Click View as XML and add the following lines in the text box:

```
<property>
 <name>hive.server2.authentication</name>
 <value>KERBEROS</value>
</property>
```

Alternatively, you can click + to enable the Editor mode and specify `hive.server2.authentication` in the Name field and `KERBEROS` in the Value field.

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

### Results

The `hive.server2.authentication` property is appended to the `hive/conf/hive-site.xml` file. From now on, Hive will use Kerberos to authenticate access requests from the Hive editor within Hue and LDAP when you access Hive using Beeline.

## 502 Proxy Error while accessing Hue from the Load Balancer

When you access Hue from the Hue Load Balancer and encounter the "502 Proxy Error Proxy Error The proxy server received an invalid response from an upstream server. The proxy server could not handle the request POST /desktop/api/search/entities." error message, then increase the proxy timeout value for the Hue Load Balancer using Cloudera Manager.

### Procedure

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration Scope Load Balancer Load Balancer Advanced Configuration Snippet (Safety Valve) for `httpd.conf`.
3. Add the following line in the Load Balancer Advanced Configuration Snippet (Safety Valve) for `httpd.conf` text box:

```
ProxyTimeout 600
```



**Note:** If you are seeing the following error, then considering increasing the proxy timeout value to 1000 seconds: Proxy Error Proxy Error The proxy server received an invalid response from an upstream server. The proxy server could not handle the request POST /notebook/api/get\_logs.

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

## Invalid method name: 'GetLog' error after submitting Hive queries

The Invalid method name: 'GetLog' (code THRIFTAPPLICATION): None error can occur after you submit Hive queries from the Hue editor and while Hue tries to fetch the result set. If you encounter this error, then set the `beeswax use_get_log_api` property to false using Cloudera Manager.

### Procedure

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration.
3. Add/update the `beeswax` section as follows in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[beeswax]
use_get_log_api=false
```

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

## "Authorization Exception" error on submitting queries in Hue

If you have secured your cluster using Ranger, then you must grant the required permissions to your users and groups from the Ranger web UI. If your users do not have proper permissions, then they may not be able to access certain databases or tables from the Hue editor.

### About this task

If your users see the "Authorization Exception: User does not have privileges to execute..." error on submitting queries from the Hue editor, then grant them the proper permissions using the Ranger web UI.

### Procedure

1. Log into Cloudera Manager as an administrator.
2. Go to Clusters Ranger service Instances tab and note down the hostname corresponding to the "Ranger Usersync" role type.
3. Open the Ranger web UI by clicking Ranger Admin Web UI.
4. SSH into the Ranger Usersync host that you noted in step 2 and add the user or the group as follows:

```
ssh root@example.domain.site useradd [***USERNAME/GROUP-NAME***] passwd
[***PASSWORD***]
```

5. On the Ranger web UI, click Hadoop SQL listed under the HADOOP SQL service.  
The **Hadoop SQL Policies** page is displayed.

6. On the **Hadoop SQL Policies** page, you can grant the new user access to all the databases or to specific databases by adding a new policy.
- To grant the permission on all databases:
    - Click the policy ID corresponding to "all - database, table, column".

List of Policies : Hadoop SQL

Search for your policy...

Add New Policy

Policy ID	Policy Name	Policy Labels	Status	Audit Logging	Roles	Groups	Users	Action
7	all - global	-	Enabled	Enabled	--	--	hive beacon dpprofiler hue + More...	
8	all - database, table, column	-	Enabled	Enabled	--	--	hive beacon dpprofiler hue + More...	
9	all - database, table	-	Enabled	Enabled	--	--	hive beacon dpprofiler hue + More...	
10	all - database	-	Enabled	Enabled	--	public	hive beacon dpprofiler hue + More...	
11	all - hiveservice	-	Enabled	Enabled	--	--	hive beacon dpprofiler hue + More...	

- On the **Edit Policy** page, add the user whom you want to grant the permission in the Select User field under the Allow Conditions section as shown in the following image:

Allow Conditions :

hide

Select Role	Select Group	Select User	Permissions	Delegate Admin	
Select Roles	Select Groups	<input type="checkbox"/> hive <input type="checkbox"/> beacon <input type="checkbox"/> dpprofiler <input type="checkbox"/> hue <input type="checkbox"/> admin <input type="checkbox"/> impala	<input checked="" type="checkbox"/> select <input checked="" type="checkbox"/> update <input checked="" type="checkbox"/> Create <input checked="" type="checkbox"/> Drop <input checked="" type="checkbox"/> Alter <input checked="" type="checkbox"/> Index <input checked="" type="checkbox"/> Lock <input checked="" type="checkbox"/> All <input checked="" type="checkbox"/> Read <input checked="" type="checkbox"/> Write <input checked="" type="checkbox"/> ReplAdmin <input checked="" type="checkbox"/> Service Admin <input checked="" type="checkbox"/> Temporary UDF Admin <input checked="" type="checkbox"/> Refresh	<input checked="" type="checkbox"/>	<input type="checkbox"/>
Select Roles	Select Groups	<input type="checkbox"/> rangerlookup	<input checked="" type="checkbox"/> Read	<input type="checkbox"/>	<input type="checkbox"/>
Select Roles	Select Groups	<input type="checkbox"/> (OWNER)	<input checked="" type="checkbox"/> All	<input checked="" type="checkbox"/>	<input type="checkbox"/>

To grant permissions to a group, enter the group name in the Select Group field.

- Click Save.
- To grant permission on specific database:
    - Click Add New Policy.

The **Create Policy** page is displayed.

- Under the Policy Details section, specify the policy name and select the database, table, and column that you want your user to access as shown in the following image:

Service Manager > Hadoop SQL Policies > Create Policy

Policy Details :

Policy Type: Access

Policy Name: default database, tables, column ☐ enabled ☐ no

Policy Label: Policy Label

database: ☐ default ☐ Include

table: ☐ \* ☐ Include

column: ☐ \* ☐ Include

Description: This policy grants users permissions to only the "default" database and its tables and columns.

Audit Logging: YES

Allow Conditions :

hide

Select Role	Select Group	Select User	Permissions	Delegate Admin	
Select Roles	Select Groups	Select Users	<input checked="" type="checkbox"/> select <input checked="" type="checkbox"/> update <input checked="" type="checkbox"/> Create <input checked="" type="checkbox"/> Drop <input checked="" type="checkbox"/> Alter <input checked="" type="checkbox"/> Index <input checked="" type="checkbox"/> Lock <input checked="" type="checkbox"/> All <input checked="" type="checkbox"/> Read <input checked="" type="checkbox"/> Write <input checked="" type="checkbox"/> ReplAdmin <input checked="" type="checkbox"/> Service Admin <input checked="" type="checkbox"/> Temporary UDF Admin <input checked="" type="checkbox"/> Refresh <input checked="" type="checkbox"/> Select/Deselect All	<input type="checkbox"/>	<input type="checkbox"/>

Add Permissions

- Under the Allow Conditions section, enter the username in the Select User field and click Add Permissions and select the permissions that your user must have.



To grant permissions to a group, enter the group name in the Select Group field.

d. Click Add.

7. Start the Hue service from Cloudera Manager.

### Results

The user or the group should be able to run any query on any entities as defined in the policy.

## Cannot alter compressed tables in Hue

Due to a known bug in the Oracle database (12c and higher), you cannot perform ALTER TABLE operations (add, delete, drop, modify) on compressed tables. If you have compressed tables in the Hue schema, then you may see the “ORA-39726: unsupported add/drop column operation on compressed tables” error.

### About this task

Even if you uncompress an existing table, you may not be allowed to alter the columns. To resolve this issue:

### Procedure

1. SSH into the host on which you have installed the Oracle database.
2. Create a new uncompressed table with the same structure as the compressed table.
3. Copy the data from the compressed table to the new uncompressed table.
4. Rename or delete the compressed table.
5. Rename the uncompressed table with the name of the original compressed table.

### Results

You should now be able to perform ALTER TABLE operations (add, delete, drop, modify) on the Hue tables.

## Connection failed error when accessing the Search app (Solr) from Hue

If you are using Solr with Hue to generate interactive dashboards and for indexing data, and if you have deployed two Solr services on your cluster and selected the second one as a dependency for Hue, then Cloudera Manager assigns the hostname of the first Solr service and the port number of the second Solr service resulting in an incorrect Solr URL in the search section of the hue.ini file. As a result, you may see a “Connection failed” error when you try to access the Search app from the Hue web UI.

### About this task

For example, consider two Solr services with the following configuration:

Solr-1, hostname=solr1, port:2345

Solr-2, hostname=solr2, port=4567

If you select Solr-2 as a dependent service for Hue, then Cloudera Manager updates the search section of the hue.ini file as follows:

```
[search]
URL of the Solr Server
solr_url=http://solr2:2345/solr/
```

As a result, you may not be able to access the Search app from the Hue web UI. To resolve this issue:

### Procedure

1. Log into Cloudera Manager as an administrator.
2. Go to **Clusters Hue service Configuration** and add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[search]
URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

```
solr_url=http://solr2:4567/solr/
```

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

## Downloading query results from Hue takes time

If downloading query results from the Hue web UI takes time or if the operation exits with the "Invalid query handle" message, then you can improve the processing speed by increasing the number of threads used by the Hue web server.

### Procedure

1. Sign in to Cloudera Manager as an Administrator.
2. Go to **Clusters Hue service Configuration** and search the `cherrypy_server_threads` property.  
The Hue Web Server Threads field is displayed.
3. Increase the thread count to 100 or a higher value.  
The default Hue web server thread count is 50.
4. Click **Save Changes**.
5. Restart the Hue service.

## Hue Load Balancer does not start after enabling TLS

The Hue Load Balancer reads the private key file that is defined in the Hue Load Balancer TLS/SSL Server Private Key File (PEM Format) configuration property to start. Because the private key files are usually encrypted, the Hue Load Balancer must be configured to use the corresponding key password, without which it cannot start.

### About this task

If you have enabled TLS for the Hue service on your cluster, and if the private key file is password protected (encrypted), then you may see the following error in the Hue Load Balancer log file (`/var/log/hue-httpd/error_log`):

```
AH02312: Fatal error initialising mod_ssl, exiting.
```

The following message is also logged in the `/var/run/cloudera-scm-agent/process/[***XXX-HUE_LOAD_BALANCER***]/logs/stdout.log` file:

```
CLOUDERA_HTTPD_USE_SSL=true
Apache/2.4.6 mod_ssl (Pass Phrase Dialog)
Some of your private key files are encrypted for security reasons.
In order to read them you have to provide the pass phrases.

Server example.test.com:443 (RSA)
```

```
Enter pass phrase:
```

To resolve this issue:

### Procedure

1. Create a password file in your chosen security directory and insert the private key password as shown in the following example:

```
echo "abc123" > /etc/security/password.txt
```

Where abc123 is the private key password and password.txt is the password file.

2. Set the file ownership and permissions as shown in the following example:

```
chown hue:hue password.txt
chmod 700 password.txt
```

3. Enter the path to the file containing the passphrase used to encrypt the private key of the Hue Load Balancer server in the Hue Load Balancer TLS/SSL Server SSLPassPhraseDialog field.

In this case, /etc/security/password.txt.

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

## Unable to terminate Hive queries from the Hue Job Browser in a Kerberized cluster

On a Kerberized cluster, if YARN does not have Kerberos authentication enabled for HTTP web consoles, then you may not be able to terminate Hive queries from the Hue Job Browser, and you may see the following error in the Hue role log `runcpserver.log` file: "The default static user cannot carry out this operation. (error 403)".

### About this task

On a Kerberized cluster, YARN must have Kerberos authentication enabled for HTTP web consoles. If authentication is not enabled, then the user or application that is trying to access YARN using a REST API is identified as the default "dr.who" user. The default user does not have permissions to access the YARN UI and terminate the running jobs. As an immediate solution, you can terminate the job from the Hue query editor or from the YARN CLI using the following command:

```
yarn application -kill [***APPLICATION-ID***]
```

To enable terminating jobs and running queries from the Hue Job Browser, enable Kerberos authentication for HTTP web consoles for YARN as follows:

### Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to **Clusters YARN Configuration** and type `enable kerberos` in the search box.
3. Select **Enable Kerberos Authentication for HTTP Web-Consoles**.
4. Click **Save Changes**.
5. Restart the YARN service.

## Unable to view or create Oozie workflows in Hue on a Knox-secured cluster

If you are unable to view Oozie workflow actions, such as "HiveServer2 Script" or "Shell Script" on the **Oozie Editor** page in Hue on a Knox-secured cluster, then check the Hue logs for the following warning: "POST /desktop/log\_js\_error HTTP/1.1" --- Referer checking failed - https://[\*\*\*FQDN\*\*\*]:[\*\*\*PORT\*\*\*]/oozie//editor/workflow/new does not match any trusted origins.. To fix this issue, add the Oozie server URL to the trusted\_origins property in the Hue Advanced Configuration Snippet.

### About this task

When you set up Knox on your CDP cluster, Knox authenticates access or requests from other services and applications. If you specify the Oozie server URL in the trusted\_origins property, Knox can check that the incoming request is from a trusted source (Oozie) and approves access, allowing you to view and create Oozie workflows from Hue.

### 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. Add the Oozie server URL in the trusted\_origins property under the desktop section as follows:

```
[desktop]
[[session]]
Comma-separated list of Oozie nodes and Oozie ports
for each Oozie instance
trusted_origins=[***OOZIE-NODE1***]:[***OOZIE-PORT1***], [***OOZIE-NODE2***]:[***OOZIE-PORT2***], ...
```

For example:

```
[desktop]
[[session]]
Comma-separated list of Oozie nodes and Oozie ports
for each Oozie instance
trusted_origins=localhost:11000
```

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

## MySQL: 1040, 'Too many connections' exception

If Hue displays the "1040, Too many connections" exception, then it is possible that the Hue backend database is overloaded and out of maximum available connections. To resolve this issue, you can increase the value of the max\_connections property for your database.

### About this task

The 1040, 'Too many connections' exception occurs on a MySQL database when it runs out of maximum available connections. If you are using the Impala engine, you may see the following error message on the Hue web interface: OperationalError at /desktop/api2/context/computes/impala("1040: too many connections"). A similar error may be displayed for Hive. The exception is also captured in the Hue server logs.

### Before you begin

The `max_connections` property defines the maximum number of connections that a MySQL instance can accept. Uncontrolled number of connections can crash the server. Following are some guidelines for tuning the value of the `max_connections` property:

- Set the value of the `max_connections` property according to the size of your cluster.
- If you have less than 50 hosts, then you can store more than one database (for example, both the Activity Monitor and Service Monitor) on the same host. If you have more than 50 hosts, then use a separate host for each database/host pair. The hosts need not be reserved exclusively for databases, but each database must be on a separate host.
- For less than 50 hosts:
  - Place each database on its own storage volume.
  - Allow 100 maximum connections for each database and then add 50 extra connections. For example, for two databases, set the maximum connections to 250. If you store five databases on one host (the databases for Cloudera Manager Server, Activity Monitor, Reports Manager, Atlas, and Hive MetaStore), then set the maximum connections to 550.

To increase the number of maximum available connections and to resolve the "1040, Too many connections" exception:

### Procedure

1. Log in to Cloudera Manager and stop the Hue service.
2. SSH in to your database instance as a root user.
3. Check the number of available connections by running the following command:

```
grep max_conn /etc/my.cnf
```

`/etc/my.cnf` is the default location of the options file (`my.cnf`).

4. Set the new value of the `max_connections` property from the MySQL shell as per the guidelines provided above. For example:

```
mysql> SET GLOBAL max_connections = 550;
```

5. Restart the Hue service.

## Unable to connect Oracle database to Hue using SCAN

For high availability purposes, you may want Hue to stay connected to any Oracle database instances running in your cluster. Single Client Access Name (SCAN) serves as a cluster alias for databases in the cluster. Currently, Cloudera Manager does not provide an option to use SCAN to connect to the Oracle database. To use SCAN, you must temporarily install a MySQL database to create a Hue service and then specify Oracle settings in the Hue advanced configuration snippet.

### About this task

It is possible that other clusters or services may be able to connect to the database using the Oracle SQL Developer. But you may see the following error when you try to add the Hue service using the Cloudera Manager Add Service wizard and specifying SCAN on the **Setup Database** page: "Able to find the Database server, but not the specified database. Please check if the database name is correct and make sure that the user can access the database."

### Procedure

1. Sign in to Cloudera Manager as an administrator.
2. Add the Hue service using a MySQL database.

3. Check whether the Hue service is added successfully by launching the web UI.
4. After the Hue service is running, go to Cloudera Manager Clusters Hue service Configuration .
5. Add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini field:

```
[desktop]
[[database]]
port=0
engine=oracle
name=[***ORACLE-SCAN***] / [***SERVICE-NAME***]
user=[***HUE-DB-USER***]
```



**Note:** Specify port=0 because the port used for the Oracle database (1521) is part of the SCAN.

6. Enter the database password in the Hue Database Password field.
7. Click Save Changes.
8. Restart the Hue service.
9. Check whether you can access your databases from the Hue web UI.
10. Uninstall the MySQL database if no longer needed.

## Increasing the maximum number of processes for Oracle database

While using Oracle as a backend database for Hue, if you face issues connecting to the Hue service after restarting the database, then it is possible that the Hue is not able to get a new database connection. The following error in the Hue logs indicates that the maximum number of connections have exhausted: "ORA-12519: TNS:no appropriate service handler found". This can be resolved by increasing the number of available processes.

After restarting the Oracle database, if you are not able to connect to the Hue service, check the Hue logs for the "ORA-12519: TNS:no appropriate service handler found" error. The Hue logs are present in the following directory:

```
opt/cloudera/parcels/CDH-[***VERSION***]/lib/hue/build/env/lib/python2.7/site-packages/Django-1.11.20-
py2.7.egg/django/db/backends/oracle/base.py
```

If you see the above error in the logs, then work with your database administrator to check whether the maximum number of processes have exceeded. If the maximum number of processes have exceeded, then you see the following error: "ORA-00020: maximum number of processes exceeded". Increase the number of processes to resolve this issue.

### How to calculate the number of database processes, transactions, and sessions?

Cloudera recommends that you allow 100 maximum connections for each service that requires a database and then add 50 extra connections. For example, for two services, set the maximum connections to 250. If you have five services that require a database on one host (the databases for Cloudera Manager Server, Activity Monitor, Reports Manager, Cloudera Navigator, and Hive metastore), set the maximum connections to 550.

From the maximum number of connections, you can determine the number of anticipated sessions using the following formula:

```
sessions = (1.1 * maximum_connections) + 5
```

For example, if a host has a database for two services, anticipate 250 maximum connections. If you anticipate a maximum of 250 connections, plan for 280 sessions.

Based on the number of sessions, you can determine the number of anticipated transactions using the following formula:

```
transactions = 1.1 * sessions
```

Continuing with the previous example, if you anticipate 280 sessions, you can plan for 308 transactions.

Work with your Oracle database administrator to apply these derived values to your system.

Using the sample values above, Oracle attributes would be set as follows:

```
alter system set processes=250;
alter system set transactions=308;
alter system set sessions=280;
```

## UnicodeDecodeError: UTF-8 codec cannot decode byte in position

You may see an error such as the following while querying a table using the Impala editor in Hue: "'utf8' codec can't decode byte 0x91 in position 6: invalid start byte". This occurs because Hue does not handle columns containing non-UTF-8 data.

To resolve this issue, contact Cloudera Support to apply the following software patches: CDPD-22129 and CDPD-27412.

## UnicodeDecodeError: 'ascii' codec cannot decode byte in position

You may see an error such as the following while downloading Impala query results in CSV format from Hue: "UnicodeDecodeError: 'ascii' codec can't decode byte 0xc2 in position 49: ordinal not in range(128)". This occurs when the query results contain special characters of non-ASCII encoding.

### About this task



**Note:** You may not face this issue if you export the query result from Hue to Excel or if you are querying the table using the Hive editor.

In CDP, Hue is compatible with Python 2.7.x, but the Tablib library for Hue has been upgraded from 0.10.x to 0.14.x, which is generally used with the Python 3 release. If you try to download Impala query results having special characters in the result set in a CSV format, then the download may fail with the ASCII unicode decode error. To fix this issue, downgrade the Tablib library to 0.12.x as follows:

### Procedure

1. SSH into the Hue server host.
2. Change directory to the following:

```
cd /opt/cloudera/parcels/CDH-7.x/lib/
```

3. Back up the hue directory.

```
cp -R hue hue_oginal
```

4. Change to the hue directory.

```
cd hue
```

5. Install the Wheel package using pip.

```
./build/env/bin/pip install wheel
```

The Wheel package is used to avoid recompiling your software during every install.

6. Install the Python Setuptools package for Hue:

```
./build/env/bin/pip install setuptools==44.1.0
```

7. Install Tablib version 0.12.1:

```
./build/env/bin/pip install tablib==0.12.1
```

8. Update the permissions of the Tablib library and other packages that were installed during the process by using the following command:

```
chmod 0755 [***TABLIB-LIBRARY-NAME***]
```



**Attention:** The Tablib library and other packages should have the drwxr-sr-x permission. Failing to have this permission can result into an “tablib module not found” error.

9. Go to Cloudera Manager and restart the Hue service.

## Fixing authentication issues between HBase and Hue

An HBase feature improvement to the Thrift Server (HBASE-19852) may cause authentication issues between HBase and Hue, and you may see the following error while accessing the HBase tables from Hue: “Failed to authenticate to HBase Thrift Server, check authentication configurations.”

### About this task

You may also see the following error in the Hue logs: “RestException: Unable to authenticate <Response [401]>”. To fix this issue, update the HBase configurations using Cloudera Manager.

### Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to **Clusters HBase service Instances** and note the hostname of the host on which the HBase Thrift Server is running.  
If multiple Thrift Servers are configured, then find the one that Hue is configured to use.
3. Go to **Configuration HBase Thrift Server Advanced Configuration Snippet (Safety Valve) for hbase-site.xml** and add the following properties:

Field	Property
Name	hbase.thrift.spnego.principal
Value	HTTP/[***HOSTNAME-FROM-STEP2***]@REALM Substitute @REALM with the actual Kerberos realm.
Name	hbase.thrift.spnego.keytab.file
Value	hbase.keytab



4. Select the following options to enable the properties:
  - Enable HBase Thrift Http Server
  - Enable HBase Thrift Proxy Users
5. Deselect the following properties:
  - Enable HBase Thrift Server Compact Protocol
  - Enable HBase Thrift Server Framed Transport
6. If you have not enabled SSL on your cluster, but if you have Kerberized your cluster, then add the following HBase configurations, without which you may encounter a NullPointerException error while starting the HBase Thrift Server:
  - a) Add the following properties in the HBase Client Advanced Configuration Snippet (Safety Valve) for hbase-site.xml field:

Field	Property
Name	hbase.thrift.ssl.enabled
Value	false

- b) Add the following properties in the HBase Thrift Server Advanced Configuration Snippet (Safety Valve) for hbase-site.xml field:

Field	Property
Name	hbase.thrift.ssl.enabled
Value	false

7. Click Save Changes.
8. Go to Clusters Hue service Configuration Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini and add the following lines:

```
[hbase]
thrift_transport=buffered
```

9. Click Save Changes.
10. Restart the HBase and Hue services to apply stale configurations.

## Hue Load Balancer does not start due to lengthy BalancerMember Route length

The Hue Load Balancer may not start if the route name in the ROLES table exceeds 64 characters. You must manually reduce the length of the route name for each Hue server in the ROLES table to be less than or equal to 64 characters, and also ensure that it is unique.

### About this task

Cloudera Manager creates the Role name in the following format, which may sometimes cause the route name to exceed 64 characters: "Service-name-Role-name-Cluster-name-Unique-hash". You may see the following error while starting the Hue Load Balancer:

```
BalancerMember Route length must be < 64 characters
+ '[' 1 '!=' 0 ']'
+ die '/var/run/cloudera-scm-agent/process/482-hue-HUE_LOAD_BALANCER/httpd.conf is invalid.'
+ echo '/var/run/cloudera-scm-agent/process/482-hue-HUE_LOAD_BALANCER/httpd.conf is invalid.'
```

```
/var/run/cloudera-scm-agent/process/482-hue-HUE_LOAD_BALANCER/httpd.conf is
invalid.
```

To resolve this issue, modify the route name column in the ROLES table in the scm database for every Hue server to be less than or equal to 64 characters, and ensure that it is unique.

### Procedure

1. SSH into the database instance as an administrator.
2. Query the ROLES table to view table content:

```
SELECT * FROM `ROLES`;
```

3. Update the values in the "NAME" column so that the route name is less than or equal to 64 characters:

```
UPDATE ROLES SET NAME=' [***HUE-ROLE-NAME***] ' WHERE ROLE_ID=[***ROLE-ID-
NUMBER***] ;
```

```
UPDATE ROLES SET NAME='hue-6c02f47dbd7e181d293c078ea293f3da' WHERE ROLE_
ID=10;
```

## Enabling access to HBase browser from Hue

If HBase impersonation is not allowed in HDFS, then you may get an API error while trying to access the HBase browser from Hue. To resolve this issue, you must allow proxy users on the Thrift gateway, and also allow all groups form all hosts in HDFS to impersonate the hbase user.

### Procedure

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters HBase service Configuration and search for the hbase.thrift.support.proxyuser property.
3. Enable the Enable HBase Thrift Proxy Users option.
4. Click Save Changes.
5. Go to Clusters HDFS service Configuration .
6. Enter the following lines in the Cluster-wide Advanced Configuration Snippet (Safety Valve) for core-site.xml field:

```
<property>
<name>hadoop.proxyuser.hbase.hosts</name>
<value>*</value>
</property>
<property>
<name>hadoop.proxyuser.hbase.groups</name>
<value>*</value>
</property>
```

7. Click Save Changes.
8. Restart the HBase and HDFS services.

## Fixing a warning related to accessing non-optimized Hue

If you are using a third-party load balancer for Hue, then you may see a warning such as "You are accessing a non-optimized Hue, please switch to one of the available addresses: https://[\*\*\*LB-HOST-FQDN\*\*\*]:[\*\*\*LB-PORT\*\*\*]". To fix this issue, you must add the URL of the Hue instance for which you are seeing the warning to the list of Hue load balancers in the Advanced Configuration Snippet.

### 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 add the following line under the [desktop] section:

```
[desktop]
hue_load_balancer=https://[***LB-HOST-FQDN***]:[***LB-PORT***]
```

You can specify a comma-separated list of URLs for which you see the warning in the hue\_load\_balancer property.

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

## Hue load balancer does not start after enabling TLS

If the Hue load balancer does not start after you enable TLS, then check whether your TLS private key file is password protected. To resolve this issue, you must configure the Hue load balancer to use the corresponding key password using Cloudera Manager.

### About this task

At startup, the Hue load balancer reads the private key file as defined in the Cloudera Manager Clusters Hue service Configuration Hue Load Balancer TLS/SSL Server Private Key File (PEM Format) configuration property. If the TLS private key file is password protected, then you may see error such as Fatal error initialising mod\_ssl, exiting and Some of your private key files are encrypted for security reasons. In order to read them you have to provide the pass phrases. in the /var/log/hue-httpd/error\_log Hue load balancer log file.

### Procedure

1. SSH into the Hue host as an Administrator.
2. Create a password file in your chosen security directory and insert the private key password, as shown in the following example:

```
echo "abc123" > /etc/security/password.txt
chown hue:hue password.txt
chmod 700 password.txt
```

3. Log in to Cloudera Manager as an Administrator.
4. Go to Clusters Hue service Configuration and enter the file path and the filename of the password file in the Hue Load Balancer TLS/SSL Server SSLPassPhraseDialog field.
5. Click Save Changes.
6. Restart the Hue service.