

Machine Learning

Cloudera AI Discovery & Exploration

Date published: 2020-07-16

Date modified: 2024-12-16

CLOUDERA

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

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Exploratory Data Science and Visualization.....	5
Prerequisites for Cloudera AI Discovery and Exploration.....	5
Starting Data Discovery and Visualization.....	6
Working with Data Discovery and Visualization.....	7
Setting up a Hive or Impala data connection manually.....	8
Setting up a Spark data connection.....	10
Setting up Amazon S3 data connection.....	11
Setting up a data connection to Cloudera Data Hub.....	11
Set up a DataHub data connection using raw code.....	13
Data connection management.....	13
Setting the workload password.....	14
Setting up a Custom Data Connection.....	14
Custom Data Connection Development.....	17
Developing and testing your first custom connection.....	17
Loading custom connections.....	22
Using data connection snippets.....	23
Managing default and backup data connections.....	25
API permissions for Projects.....	25
Troubleshooting: 401 Unauthorized.....	26

Troubleshooting: 401 Unauthorized when accessing Hive or Impala virtual warehouses.....	27
Troubleshooting: Existing connection name.....	27
Troubleshooting: Empty data page.....	28
Troubleshooting: Some data connections not shown.....	28

Exploratory Data Science and Visualization

Exploratory Data Science and Visualization makes it simple for data scientists to get started on a data science project.

When you start a data science project, you are presented with a blank notebook, and no indication of what data sources are available on the Cloudera platform, or how to access them. The Exploratory Data Science and Visualization feature automatically discovers the data sources available to you, from within the standard Cloudera AI user interface.

The Exploration Data Science and Visualization experience enables all features for Exploratory Data Analysis. From the Data tab you can:

- Connect to data sources that are available in your project
- Explore the data with SQL to understand its basic shape and characteristics
- Create named datasets that you can reference later
- Quickly create visualizations of the data to understand its properties
- Create dashboards that you can share with your team

For more information on what you can do in the Data tab, see the documentation for *Cloudera Data Visualization*.

Related Information

[Cloudera Data Visualization](#)

Prerequisites for Cloudera AI Discovery and Exploration

Cloudera AI Discovery and Exploration has a few prerequisites.

- For both Hive and Impala connections, the HADOOP_USERNAME and WORKLOAD_PASSWORD must be set. The HADOOP_USERNAME is set automatically by the environment.
 - Go to your **User Settings** page, select the **Hadoop Authentication** tab and log in.

Figure 1: Hadoop authentication

As a data scientist, you must set the WORKLOAD_PASSWORD.

For more information, see *Setting the workload password*.

- To manually create Impala or Hive data connections, you must have the related JDBC URL. You can obtain this from the option menu for the virtual warehouse to which you want to connect. These warehouses must be already created in the environment. For more information on how to create a data warehouse, see *Adding a New Virtual Warehouse*.
- To manually create a Spark connection, you need the Data Lake S3 bucket. This can be retrieved from the Data Lake page. For Spark data connections, you must use ML Runtimes, and specifically the Spark Runtime Add-on must be enabled before starting a workload (job or session).
- For Spark data connections, you must have permissions set correctly for external access to the S3 bucket.
- Hive and Impala data connections also require ML Runtimes. Legacy engines may work, but are not supported.
- There must be auto-discovered data connections, or your Administrator must create data connections to virtual warehouses or other data sources in your Cloudera environment.

Known Issue

- Only Hive virtual warehouses with SSO disabled are supported.

Related Information

[Create a Dashboard](#)

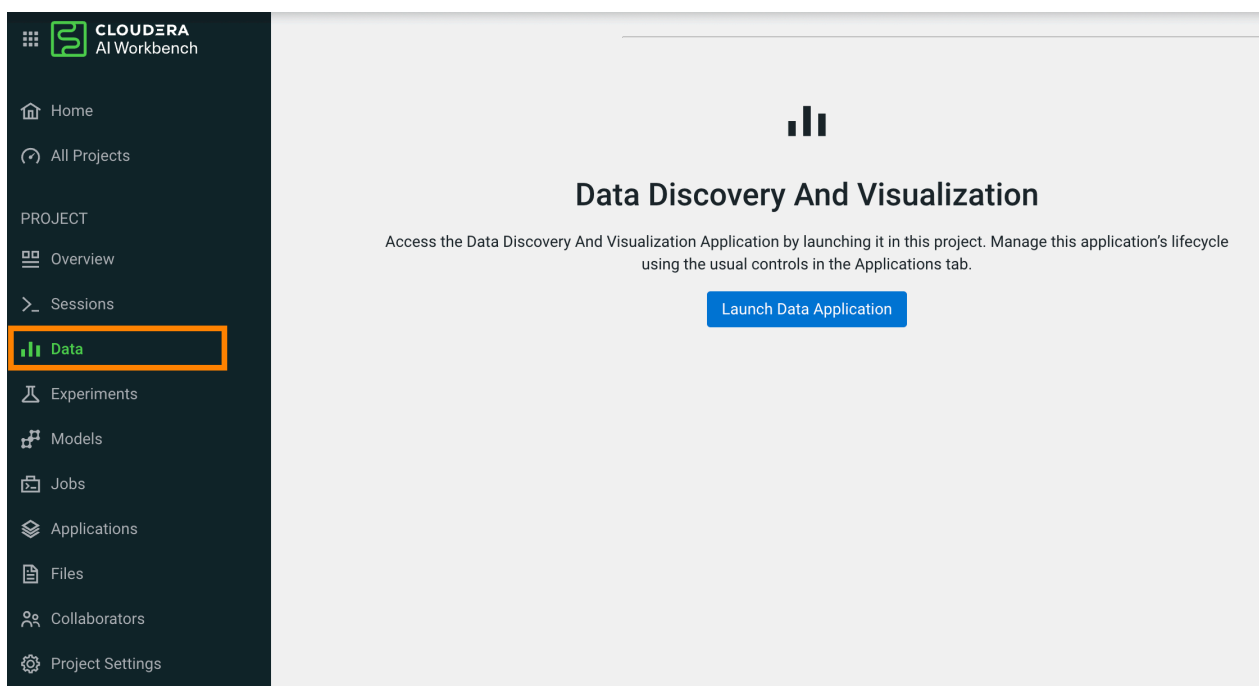
[Adding a New Virtual Warehouse](#)

[Setting the workload password](#)


Starting Data Discovery and Visualization

You can start the Exploratory Data Science and Visualization application within Cloudera AI.

1. Create or open a project.
2. In the project, click Data in the left navigation pane.
3. If the application was previously started, it appears in the UI.
4. If it is not displayed, click Launch Data Application to start the application. It takes a few minutes to start for the first time.



When the application starts, you can see four tabs at the top of the UI. From here, you can create new dashboard or dataset or try some of the tasks described in the next topics.



The screenshot displays the Cloudera AI Workbench interface. At the top, there are navigation tabs for HOME, SQL, VISUALS, and DATASETS. A sidebar on the left contains a menu with options like Home, All Projects, PROJECT (Overview, Sessions, Data, Experiments, Models, Jobs, Applications, Files, Collaborators, Project Settings), AMPs, Runtime Catalog, Learning Hub, User Settings, Site Administration, and Help. The main content area features a dashboard with several summary cards: 0 QUERIES, 16 DASHBOARDS, 1 APPS, 11 DATASETS, and 2 CONNECTIONS. Below these are sections for Recent Queries, Recent Dashboards and Apps, Recent Datasets, and Recent Connections. The Recent Connections section shows two connections: 'test-cdv-auth' (TYPE: Impala) and 'samples' (TYPE: SQLite). At the bottom, there is a 'Sample Dashboards' section with seven thumbnails: 'Taxi rides application', 'Iris species w/ images', 'Consumer View', 'Inspector View', 'Time & Industry Threat ...', 'Global Threats', and 'Census Dashboard'.

Related Information

[Creating a dataset](#)

[Create a Dashboard](#)

[Troubleshooting: 401 Unauthorized](#)

Working with Data Discovery and Visualization

This section provides examples for tasks you can perform using Data Discovery and Visualization.

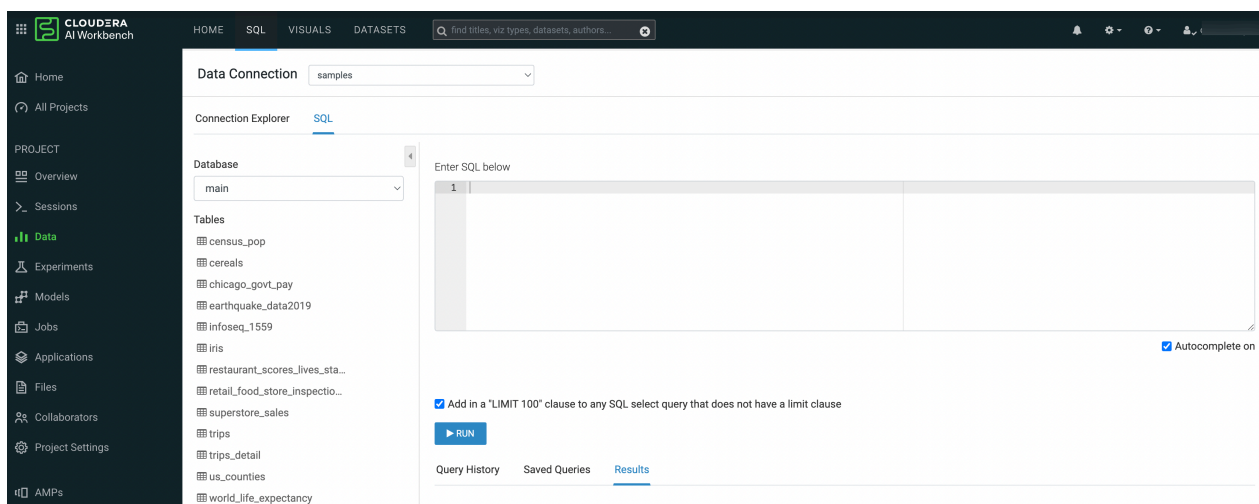
Run an SQL query

On the **SQL** tab, select a Data Connection from the dropdown box.



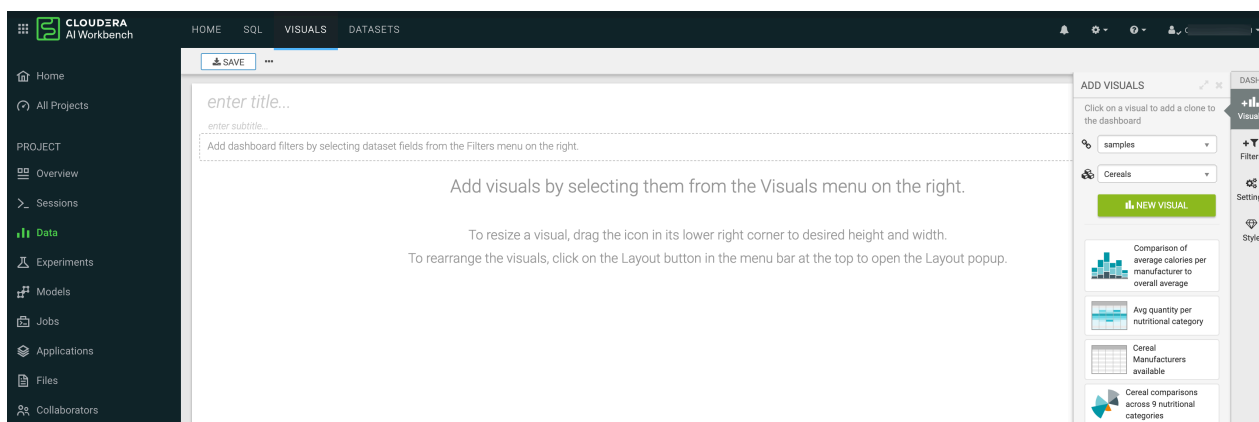
Note: If you do not see the data connections that you have configured in the Project, refer to the *Troubleshooting* details for more information.

Once you select a Data Connection, you can explore the available databases and tables. You can start writing an SQL query and click the table names to insert them into the query string. For more information, see *Creating a dataset*.



Build a dashboard

Click New Dashboard to access the Dashboard Designer. You can select from a wide variety of graphs, click to add them to your dashboard, and select the columns from your dataset to use in the visual. For more information, see *Create a Dashboard*.



Related Information

[Creating a dataset](#)

[Create a Dashboard](#)

[Troubleshooting: 401 Unauthorized](#)

Setting up a Hive or Impala data connection manually

Data connections to Hive or Impala virtual warehouses within the same environment as the Cloudera AI Workbench are automatically discovered and configured. You can also set up a data connection manually, which works across Cloudera environments. Follow this procedure to set up a Hive or Impala data connection.

Procedure

1. Log into the Cloudera web interface and navigate to the Cloudera Data Warehouse service.
2. In the Cloudera Data Warehouse service, select Virtual Warehouses in the left navigation panel.

3. Select the options menu for the warehouse you want to access, and select Copy JDBC URL.

Virtual Warehouses

cdw-hive-2
Stopped
compute-1630362293-g4vx
eng-ml-dev-env-aws-dl-default

NODE COUNT	TOTAL CORES	TOTAL MEMORY	TYPE
0	12	56 GB	HIVE

imesz-hive-test
Stopped
compute-1630331041-8wx5
imesz-test-huebb

NODE COUNT	TOTAL CORES	TOTAL MEMORY	TYPE
0	12	56 GB	HIVE UNIFIED

imesz-impala-test
Stopped
impala-1630330335-sjgz
imesz-test-huebb

- Clone
- Edit
- Delete
- Upgrade
- Copy JDBC URL**
- Download JDBC Jar
- Open DAS
- Open HUE
- Set Compactor
- Run AutoScaling Demo
- Collect Diagnostic Bundle


4. Return to the Cloudera AI service. In Site Administration Data Connections , select New Connection.
5. Enter the connection name. You cannot have duplicate names for data connections within a workbench or within a given project.
6. Select the connection type:
 - a. Hive Virtual Warehouse
 - b. Impala Virtual Warehouse
7. Paste the JDBC URL for the data connection.

8. (Optional) Enter the Virtual Warehouse Name. This is the name of the warehouse in Cloudera Data Warehouse.

New Data Connection ✕

*** Name**

*** Type** ⓘ

 Hive Virtual Warehouse ▾

*** JDBC URL** ⓘ

Virtual Warehouse Name ⓘ

Available ⓘ

Cancel Create

Results

The data connection is available to users by default. To change availability, click the Available switch. This switch determines if the data connection is displayed in Projects created within the workbench.

Setting up a Spark data connection

Spark data connections within the same environment as Cloudera AI are automatically discovered, but you can also set up a connection manually. Follow this procedure to set up a Spark data connection.

Procedure

1. In the Workbench UI, select the link environment for the workbench you are using. This takes you to the Environments UI.

2. In Environments, select **Data Lake Cloud Storage** tabs.
3. Select the directory path shown for Hive Metastore External Warehouse, and copy it.
4. In Project Settings > Data Connections, click **New Connection**.
5. Enter a name for the connection.
6. Select the type: **Spark Data Lake**
7. Paste the value you copied in step 3 into **Datalake Hive Metastore External Warehouse Directory**.
8. Click **Create**.

Results

The data connection is available to users by default. To change availability, click the **Available** toggle.

Setting up Amazon S3 data connection

Amazon S3 object store connection is automatically created for Cloudera AI Workbenches to make it easier to connect to the data stored within the same environment. Other Data Connections can be configured to other S3 locations manually. You can also set up a connection manually.

About this task

Amazon S3 data connections are available only on AWS workbenches where RAZ (Ranger Authorization Service) is enabled to authorize connections to the environment's S3 buckets.

Procedure

1. In the **CDP** console, click the **Cloudera AI** tile.
The **Home** page displays.
2. Click **Site Administration** in the left navigation menu.
The **Site Administration** page displays.
3. In the Data Connections page, click **New Connection**.
The **New Data Connection** window is displayed.
4. In the Name field, enter a name for the connection.
5. In the Type drop-down list, select the type as **S3 Object Store**.
6. Click **Create**.

This data connection only supports per-bucket operations. For information on using the S3 data connection, and connection wrapper for the S3 boto client, see *Amazon Boto documentation*.

The data connection is available to users by default. To change availability, click the **Available** toggle. This switch determines if the data connection is displayed in Projects created within the workbench.

Related Information

[Amazon Boto documentation](#)

Setting up a data connection to Cloudera Data Hub

You can set up a data connection to the DataHub cluster. You can set up a connection using the **New Connection** dialog, or by using raw code inside your project. Both approaches are shown below.

Procedure


1. Log into the console web UI.

2. Depending on which connection you want to use, click on either **Connect to Hive** or **Connect to Impala** tile.
3. Ensure your Data Hub cluster name is correct in the popup.
4. Copy the JDBC URL string.
5. Now click on the **Build a Data Science Project** tile to log into the the Cloudera AI Workbench.
6. In **Site Administration Data Connections** , select **New Connection**.
7. Return to the Cloudera AI service. In **Site Administration Data Connections** , select **New Connection**.
8. Enter the connection name. You cannot have duplicate names for data connections within a workbench or within a given project.
9. Select the connection type:
 - a. **Hive Virtual Warehouse**
 - b. **Impala Virtual Warehouse**
10. Paste the JDBC URL for the data connection.
11. (Optional) Enter the Virtual Warehouse Name. This is the name of the warehouse in Cloudera Data Warehouse.

New Data Connection ✕

*** Name**

*** Type** ⓘ

 Hive Virtual Warehouse ▼

*** JDBC URL** ⓘ

Virtual Warehouse Name ⓘ

Available ⓘ

Results

The data connection is available to users by default. To change availability, click the Available switch. This switch determines if the data connection is displayed in Projects created within the workbench.

Set up a DataHub data connection using raw code

It is recommended to use the New Connection dialog to create a new data connection. If needed, you can also set up a data connection in your project code by using and adapting the following code snippet.

Example

```
from impala.dbapi import connect

#Example connection string:
# jdbc:hive2://my-test-master0.eng-ml-i.svbr-nqvp.int.cldr.work/;ssl=true;
transportMode=http;httpPath=my-test/cdp-proxy-api/hive

conn = connect(
    host = "my-test-master0.eng-ml-i.svbr-nqvp.int.cldr.work",
    port = 443,
    auth_mechanism = "LDAP",
    use_ssl = True,
    use_http_transport = True,
    http_path = "my-test/cdp-proxy-api/hive",
    user = "csso_me",
    password = "Test@123")
cursor = conn.cursor()
cursor.execute("select * from 3yearpop")

for row in cursor:
    print(row)
cursor.close()
conn.close()
```

Data connection management

Consider the following guidelines about data connection management.

- Manage data connections in a workbench

At the workbench level, you can check the data connections that are available in a workbench. In Project Settings Data Connections, check that your desired data source is present. You can also set the availability for any discovered connections, if necessary.

- Manage data connections in a project

All the data connections that were available in the workbench when the project was created will be automatically created in the project as well. The available connections can then further be marked as unavailable if desired so. You can update any changes to the connections that were made at the workbench level by clicking Sync with Workbench. Any changes made here only apply to your project.

- Data connection availability

Consider these two scenarios for setting data connection availability.

1. If a workbench data connection is marked Unavailable, and you then create a project, the data connection will not appear in the project. If the connection is then changed to Available, and then the Sync with Workbench button is clicked, the connection will appear in the project.
2. If a workbench data connection is marked Available, and you then create a project, the connection shows up. If the workbench data connection is then toggled to Unavailable, and you click Sync with Workbench in the project, the data connection will remain available in the project.

Setting the workload password

Set the workload password to connect your code to the data.

You need to set an environmental variable for the password (`WORKLOAD_PASSWORD`). This environmental variable is needed to access a Hive or Impala warehouse. The password you set in Cloudera AI must match the password you set in Management Console. For more information on setting the workload password in Management Console, see *Setting the workload password*.

To set the password:

1. In Cloudera AI, go to User Settings Environment .
2. Under the Reserved Environment Variables set your `WORKLOAD_PASSWORD`.
3. Click Save.
4. In Management Console, select your user name in the left navigation pane, and select Profile.
5. Click Set Workload Password.
6. Enter the same password as in step 2, and click Set Workload Password.



Note: You need to restart any sessions or applications that are running in your workbench.

The screenshot shows the Cloudera AI Workbench interface. The main content area is titled 'User Settings' and has a sub-tab for 'Environment Variables'. Under 'Reserved Environment Variables', there is a table with one row: 'WORKLOAD_PASSWORD' with the value 'WORKLOAD_PASSWORD' and a 'Save' button. Below this is the 'Environment Variables' section, which includes a description and an 'Add Variable' button. The left sidebar contains a navigation menu with items like Home, Projects, Sessions, Experiments, Models, Jobs, Applications, AMPs, Runtime Catalog, Learning Hub, User Settings, and Site Administration.

Related Information

[Setting the workload password](#)

Setting up a Custom Data Connection

Data connections that point to data sources outside of Cloudera or require custom configurations can be created and made available to end users with Custom Data Connections. These Python implementations of the Cloudera AI Data

library are stored in the Data Connections Registry. Workbench users can track and connect to any data source and connection implementation a Cloudera AI Administrator makes available.

About this task

Consider the followings:

- Custom connections can only be created in projects created by the Administrator.
- The project source selection list in the Data Connection creation dialogue only displays projects created by the user.
- Team projects or projects with multiple collaborators will also not be displayed, only those directly created by the user.
- Custom connections at workbench level can only be edited by the creator, not other Administrator users. Attempts at editing workbench level custom connections will result in an error.

Before you begin

Before setting up a custom connection, you might want to create a dedicated Cloudera AI Team to collaborate on external connections. A good practice is to separate the connection code projects and and configure collaborators on the Team level to build and maintain the connection code.

Procedure

1. Develop your own custom data connection (see *Developing a Custom Data Connection*) in a Cloudera AI project, or clone an existing custom data connection files directory into a Cloudera AI project.
2. In Site Administration Data Connections , select New Connection.
3. Enter the connection name. You cannot have duplicate names for data connections within a workbench or within a given project.
4. Select the connection type: Custom Connection
5. Enter the Type Display name. This should be a descriptive label to help Cloudera AI project owners identify what this custom connection could be used for.
6. Select the Cloudera AI Project and Project directory which contains your custom connection implementation
 - a. Connection files must be in a directory and not in the root of your project.
 - b. A snapshot of all implementation files in the directory will be uploaded to the Cloudera AI Custom Data Connection registry located in the workbench.
 - c. These uploaded files are safe from any changes to the originating project. To make changes to the files, create a new custom data connection.

- (Optional) Enter any custom parameters. These are available during a session and can be validated or overridden depending on the interface implementation for the custom data connection. Refer to the implementation of your custom data connection for specific details on required keys and values.

New Data Connection



*** Name**

*** Type** ⓘ

*** Type Display**

*** Project** ⓘ

Custom Parameters ⓘ

- +

- +

Available ⓘ

- Click Create.

Results

The data connection is now available to all users. To change availability, click the Available switch. This switch determines if the data connection is displayed in Projects created within the workbench. Refer to *Data connection management* for availability of your newly created custom connections in new and existing Cloudera AI Projects.

Related Information

[Data connection management](#)

Custom Data Connection Development

Custom data connections can be developed from within Cloudera AI Workbench and Python Sessions using the Cloudera AI Python Data Library and implementing the Cloudera AI Custom Connection Interface.

You can view CustomConnection interface help descriptions within in a session:

```
import cml.data_v1 as cmldata
help(cmldata.customconnection)
```

Alternatively, you can inspect the source content as follows:

```
import cml.data_v1 as cmldata
import inspect
print(inspect.getsource(cmldata.customconnection))
```

Your custom connection code must implement the CustomConnection interface for the cml.data_v1 library to load your module dynamically (see *Loading custom connections*)

Two functions are already implemented so that the Cloudera AI Python Data Library can dynamically load your Python module implementation and make custom parameters available in self.parameters. In most cases, you will not need to reimplement these:

1. `__init__(self, properties)`
2. `update_properties(self, properties)`

The rest of the interface functions are included as common functions that you may want to implement.

1. `get_base_connection(self)`
2. `get_pandas_dataframe(self, query)`
3. `get_cursor(self)`
4. `print_usage(self)`
5. `override_parameters(self)`

See *Developing and testing your first custom connection* for a simple example of how to implement these.

Related Information

[Loading custom connections](#)


[Developing and testing your first custom connection](#)

Developing and testing your first custom connection

You can develop a custom connection in your own project, as shown in this example.

Procedure

1. Create a new directory `hello-world-conn-dev` and file `hello-conn.py` to contain your custom data connection files.



```

File Edit View Navigate Run ▶ hello-world-conn-dev/hello-conn.py
1 """Hello World Dataconn Connection Implementation"""
2 from cml.data_v1.customconnection import CustomConnection
3
4
5 class HelloWorld(CustomConnection):
6
7     """ Implementation of the interface function print_usage"""
8     def print_usage(self):
9         print(""" This is a %s connection""% (self.app_name))
10
11     """ Implementation of custom function for hello connections"""
12     def hello_world(self):
13         print("These are the specified parameters" )
14         print (self.parameters)
15
16
17

```

2. Implement your custom data connection in `hello-conn.py`. (There must only be one class which implements `CustomConnection` in this directory.)

```

"""Hello World Custom Connection Implementation"""
from cml.data_v1.customconnection import CustomConnection

class HelloWorld(CustomConnection):

    """ Implementation of the interface function print_usage"""
    def print_usage(self):
        print(""" This is a %s connection""% (self.app_name))

    """ Implementation of custom function for hello connections"""
    def hello_world(self):
        print("These are the specified parameters" )
        print (self.parameters)

```

3. Test your custom data connection locally (for more information see *Loading custom connections*)

```

test_params = {"PARAM_1": "foo", "PARAM_2": "bar"}
import cml.data_v1 as cmldata

conn = cmldata.get_custom_connection_from_local(
    "hello-world-conn-dev",
    "my-hello-connection",
    test_params)
conn.print_usage()

```

```
conn.hello_world()
```

← Project >_ Terminal Access ☰ Data ✍ Clear ⚡ Interrupt ■ Stop Sessions ▾

Untitled Session  **Running**

By CDEP CREATED ACCOUNT – Session – 2 vCPU / 4 GiB Memory – a few seconds ago

Session Logs [Collapse](#) [Share](#) [Export PDF](#)

```
> test_params = {"PARAM_1": "foo", "PARAM_2": "bar"}
> conn = cmldata.get_custom_connection_from_local(
    "hello-world-conn-dev",
    "my-hello-connection",
    test_params)
| Loaded custom connection my-hello-connection
> conn.print_usage()
| This is a my-hello-connection connection
> conn.hello_world()
| These are the specified parameters
| {'PARAM_1': 'foo', 'PARAM_2': 'bar'}
```

```
>
```

4. In Site Administration Data Connections , select New Connection, and fill in the fields as shown below.

New Data Connection ✕

*** Name**

*** Type** ⓘ

*** Type Display**

*** Project** ⓘ

*** Connection Files**

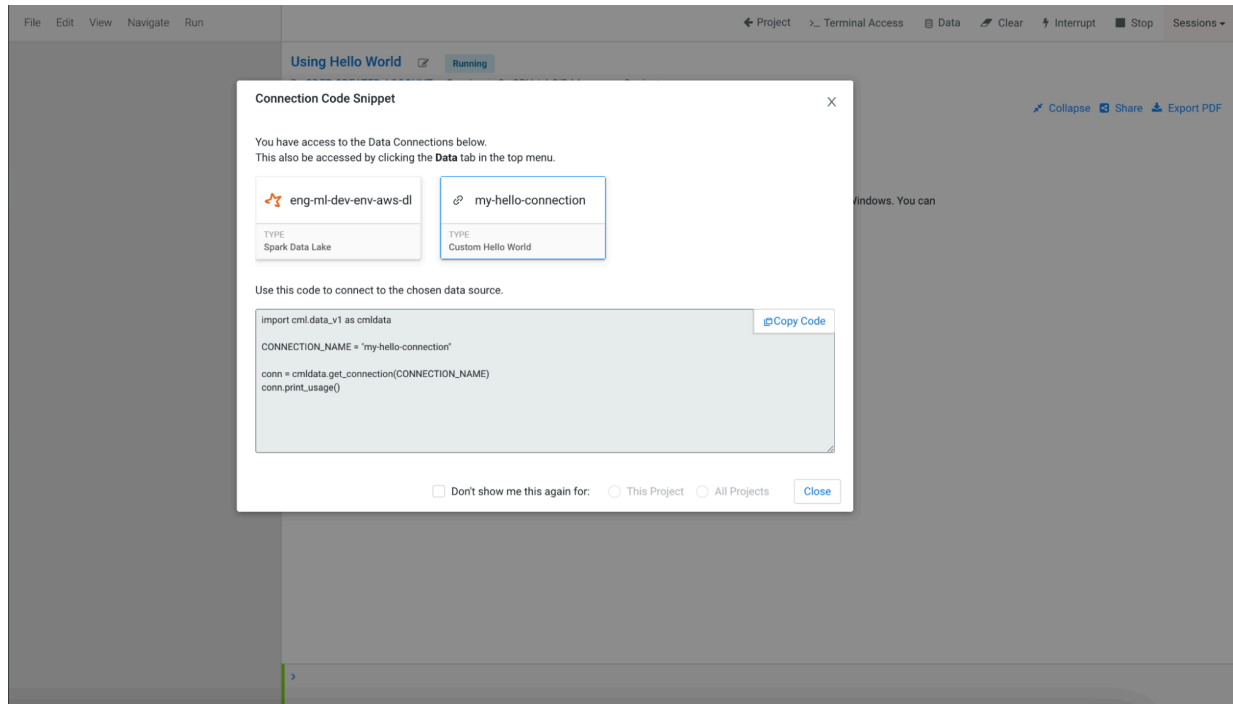
Custom Parameters ⓘ

<input type="text" value="PARAM_1"/>	<input type="text" value="foo"/>	-	+
<input type="text" value="PARAM_2"/>	<input type="text" value="bar"/>	-	+


Available ⓘ

5. See *Data connection management* to set the availability of your newly created custom connection in new and existing Cloudera AI projects.


6. Create a project, and start a session to use your completed Hello World connection.






Note: A default custom connection code snippet is automatically displayed to the user.



```
< Project >_ Terminal Access Data Clear Interrupt Stop Sessions ▾

Using Hello World  Running
By CDEP CREATED ACCOUNT — Session — 2 vCPU / 4 GiB Memory — 2 minutes ago

Session Logs  Collapse  Share  Export PDF

> import cml.data_v1 as cmldata
> CONNECTION_NAME = "my-hello-connection"
> conn = cmldata.get_connection(CONNECTION_NAME)
Loaded custom connection my-hello-connection
> conn.print_usage()
This is a my-hello-connection connection

>
```



Note: When exporting custom connection files, git is recommended. Alternatively, download the custom connection directory only, instead of all project files.

Related Information

[Loading custom connections](#)

[Data connection management](#)

[Using data connection snippets](#)

Loading custom connections

You can instantiate a local connection for testing, using the name of your custom connection directory, a sample connection name, and an optional dictionary of parameters. This local connection object can then be used to test and implement functions in your custom connection.

The following code sample loads custom connection package directories in the same way that the Cloudera AI Data Library imports a registered custom connection when called with `get_connection`:

```
get_custom_connection_from_local(package_name, connection_name, parameters={
})
```

Returns a Cloudera AI Custom Data Connection object. For testing in-development Data Connection code.

Parameters:

- `package_name` (str): The accessible package name containing custom connection code to load.
- `connection_name` (str): The connection name to be used in Custom Connection loading.
- `parameters` (dict of str: str): Mapping of custom parameter keys and values that will be loaded by the custom connection code.

Return:

A custom connection object that implements `cml.data_v1.customconnection.CustomConnection`

Usage: `conn = load_custom_connection_source("myconndir", {"HOSTNAME": "my.instance.host.com"})`

Note: When you make changes to your custom connection file, `get_custom_connection_from_local` will dynamically re-import the contents, so the latest code on disk is instantiated for the local connection.

To load any created data connection that is available in the Cloudera AI project, use the `get_connection` instruction.

```
get_connection(dataconnection_name, parameters=None)
```

Usage: `conn = get_connection(connection_name)`

If the specified connection is of type “Custom” a snapshot of the Custom Data Connection files specified at the time of Connection creation is imported.

Using data connection snippets

As a data scientist, you can connect your project to data with a data connection snippet.

Procedure

1. Select New Project.
2. Enter the project name.
3. Start a new session. To use a Spark connection, make sure to select an ML Runtimes engine.




- In the Connection Code Snippet pane, select the data connection to use.

Connection Code Snippet



You have access to the Data Connections below.

This popup can also be accessed by clicking the **Data** tab in the top menu.

 eng-ml-dev-env-aws-dl TYPE Spark Data Lake	 eng-ml-hive-no-sso TYPE Hive Virtual Warehouse	 eng-ml-impala TYPE Impala Virtual Warehouse
--	--	---

Use this code to connect to the chosen data source.

```
import cml.data_v1 as cmldata

CONNECTION_NAME = "eng-ml-dev-env-aws-dl"
conn = cmldata.get_connection(CONNECTION_NAME)
spark = conn.get_spark_session()

# Sample usage to run query through spark
EXAMPLE_SQL_QUERY = "show databases"
spark.sql(EXAMPLE_SQL_QUERY).show()
```


[Copy Code](#)

Don't show me this again for: This Project All Projects

[Close](#)

- Select Copy Code.
- Paste the snippet into your code.

Figure 2: Pasting snippet into your code



The screenshot shows a code editor with a file named `test.py`. The code in the editor is:

```
1 import cml.data_v1 as cmldata
2
3 CONNECTION_NAME = "eng-ml-dev-env-aws-dl"
4 conn = cmldata.get_connection(CONNECTION_NAME)
5 spark = conn.get_spark_session()
6
7 # Sample usage to run query through spark
8 EXAMPLE_SQL_QUERY = "show databases"
9 spark.sql(EXAMPLE_SQL_QUERY).show()
10
```

The terminal output shows the execution of the code:

```
Test data connection [Running]
By [User] - Session - 2 vCPU / 4 GiB Memory - 17 minutes ago
Session Logs Spark UI [Collapse] [Share] [Export PDF]
> import cml.data_v1 as cmldata
> CONNECTION_NAME = "eng-ml-dev-env-aws-dl"
> conn = cmldata.get_connection(CONNECTION_NAME)
> spark = conn.get_spark_session()

Setting spark.hadoop.yarn.resourcemanager.principal to csso_wclemens
Setting spark.hadoop.yarn.resourcemanager.principal to csso_wclemens

Sample usage to run query through spark
> EXAMPLE_SQL_QUERY = "show databases"
> spark.sql(EXAMPLE_SQL_QUERY).show()

Hive Session ID = 42cabd33-d1ac-4599-846c-c9831db3e2cc
Hive Session ID = 42cabd33-d1ac-4599-846c-c9831db3e2cc
+-----+
|   databaseName|
+-----+
|         default|
|information_schema|
|              sys|
+-----+
```

- Edit the username and password, if necessary. (This is not necessary for the Spark Data Lake connection, or if the `WORKLOAD_PASSWORD` environmental variable is set up for the user.)
- Enter the data connection name.

9. Uncomment and edit the SQL query.

10. Select Run.

Results

The results of the SQL query display in the command window.

What to do next

When you have finished exploring the available data sources, you can select **Don't show me this again** for this project on the Connection Code Snippet pane, and it will no longer display when you start a session in the project.

Managing default and backup data connections

As a data scientist, you may want to set default and backup data connections (virtual warehouses). This makes it easy to manage the case where the default data source (virtual warehouse or data lake) becomes unavailable, for example.

Before you begin

You need Administrator privileges to perform this task. Here, the data connection names of **default** and **backup** are used as examples.

Procedure

1. In Site Administration Data connections, click Edit to change the name of the default data connection to default.
2. Change another connection to backup.
3. In Project Settings, click Sync with Workbench to update the names of the data connections. If two different data connections have the same name, an error occurs during synchronization.
4. If the default data connection becomes unavailable, the Workbench Administrator can go to the Data Connections tab, and rename the connections.

For example, after changing default to unavailable, change backup to default. Projects that use library code for their connection to the default data connection will continue to operate, because it is now using the new default connection.



Note: Select Sync with Workbench at the project level to make the updated connections available.

API permissions for Projects

The tasks available to collaborators depends on their level of access.

A collaborator with read-only access is able to do the following:

- List data connections
- Get data connection details



Note: The list of data connections is unavailable in the UI to viewer-only users because the Setting tab in a project is hidden for viewer collaborators. However, viewer-only users can use curl to submit API requests to list the data connections.

A collaborator with write access can:

- Create data connections
- Edit data connections
- Delete data connections

- List data connections
- Get data connection details

Troubleshooting: 401 Unauthorized

Troubleshoot the session returning a 401: Unauthorized error message.

When you are in a session and try to run the code for a Hive or Impala connection, the session returns a 401: Unauthorized HTTP error message.

Figure 3: 401: Unauthorized HTTP error message

```

modify next 2 lines to update your credentials

> USERNAME = os.getenv('HADOOP_USER_NAME')
> PASSWORD = os.getenv('WORKLOAD_PASSWORD')
> conn = cmldata.getConnection({
    'CONNECTION_NAME': CONNECTION_NAME,
    'USERNAME': USERNAME,
    'PASSWORD': PASSWORD
})
⚠️ HttpError: HTTP code 401: Unauthorized
⚠️ HttpError                                Traceback (most recent call last)
<ipython-input-1-a69c5e2af83b> in <module>
      2     'CONNECTION_NAME': CONNECTION_NAME,
      3     'USERNAME': USERNAME,
----> 4     'PASSWORD': PASSWORD
      5     })

~/local/lib/python3.7/site-packages/cmldata.py in getConnection(properties)
    105     fmt_codesnippet = _getConnectionSnippet(properties)
    106     scope = {}
--> 107     exec(fmt_codesnippet, scope)
    108     return scope["conn"]

<string> in <module>

<string> in getCursors(self)

/usr/local/lib/python3.7/site-packages/impala/hiveserver2.py in cursor(self, user, configuration, convert_type
    127     log.debug('cursor(): getting new session handle')
> |

```

Solution: You need to set your workload password. See *Setting the workload password* for details. After setting the workload password, start a new session. If this error occurs in the Data tab, then restart the Data Discovery and Exploration app.

Solution: Go to the Hadoop authentication page:

Figure 4: Access Hadoop authentication page

Set your workload password. After setting the workload password, start a new session. If this error occurs in the Data tab, then restart the Data Discovery and Exploration application.

Related Information

[Setting the workload password](#)

Troubleshooting: 401 Unauthorized when accessing Hive or Impala virtual warehouses

Troubleshoot the session returning a 401: Unauthorized error message when you are accessing a Hive database or Impala virtual warehouses.

Check that the Hive or Impala Virtual Warehouse is working.

Solution: Follow these steps to check that Hive or Impala data warehouse is running, and restart it if needed.

1. In the control plane, go to Cloudera Data Warehouse.
2. Click Virtual Warehouses, and select the data warehouse for the connection.
3. Check the status of that data warehouse, and make sure that it is running or in a good state.

Troubleshooting: Existing connection name

Troubleshoot receiving the error message that the CRN or the name is a duplicate, when attempting to synchronize data connections.

Project Settings

Options Runtime/Engine Advanced SSH Tunnels Data Connections Delete Project

⚠️ Unable to sync some connections since they have the same name as existing connections. Please rename the connections listed below, then click on "Sync with workspace" again:

- test

Show Available Only

Availability in Project	Connection Name	Connection Type	Virtual Warehouse Name	Created At	Actions
<input checked="" type="checkbox"/> Available	test	Hive Virtual Warehouse		08/30/2021 2:14 PM	

< 1 >

Solution: This indicates a project connection (one that is not copied from the workbench) has the same name or CRN as a workbench connection. To resolve this, you need to change the name or the CRN of the data connection at the project level.

Troubleshooting: Empty data page

Troubleshoot receiving an empty data page with non-ending spinner after launching the application.

- On the applications tab, click **Data Discovery and Visualization application**, and check the logs. Check if another user stopped or restarted the application.
- On the applications tab, click on **Settings Resource Profile**, and check if the available resources are sufficient. You can increase the resources if necessary, then restart the application.

Troubleshooting: Some data connections not shown

Troubleshoot the data application not showing all data connections.

1. See the error message to fix the workload password.
2. Set the workload password. Follow the procedure in *Setting the workload password*.
3. Restart the data application after setting the workload password.

Alternatively, you may need to select **Sync** under **Data Connections** to get all the data connections available at the project level.

Related Information

[Setting the workload password](#)