

1.19.

Creating and managing Cloudera Data Engineering Sessions

Date published: 2023-04-05

Date modified: 2023-04-13

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

Creating Sessions in Cloudera Data Engineering.....	4
Interacting with a Session in Cloudera Data Engineering.....	5
Connecting Sessions with the CDE CLI.....	6
External IDE connectivity through Spark Connect-based sessions	
(Technical Preview).....	6
Configuring external IDE Spark Connect sessions.....	7
Sample code to connect to an external IDE Spark Connect session.....	10
Troubleshooting errors when working with an external IDE Spark Connect session.....	12
Viewing logs for Cloudera Data Engineering Sessions.....	12
Viewing Spark UI for Cloudera Data Engineering Sessions.....	13

Creating Sessions in Cloudera Data Engineering

A Cloudera Data Engineering Session is an interactive short-lived development environment for running Spark commands to help you iterate upon and build your Spark workloads.

About this task

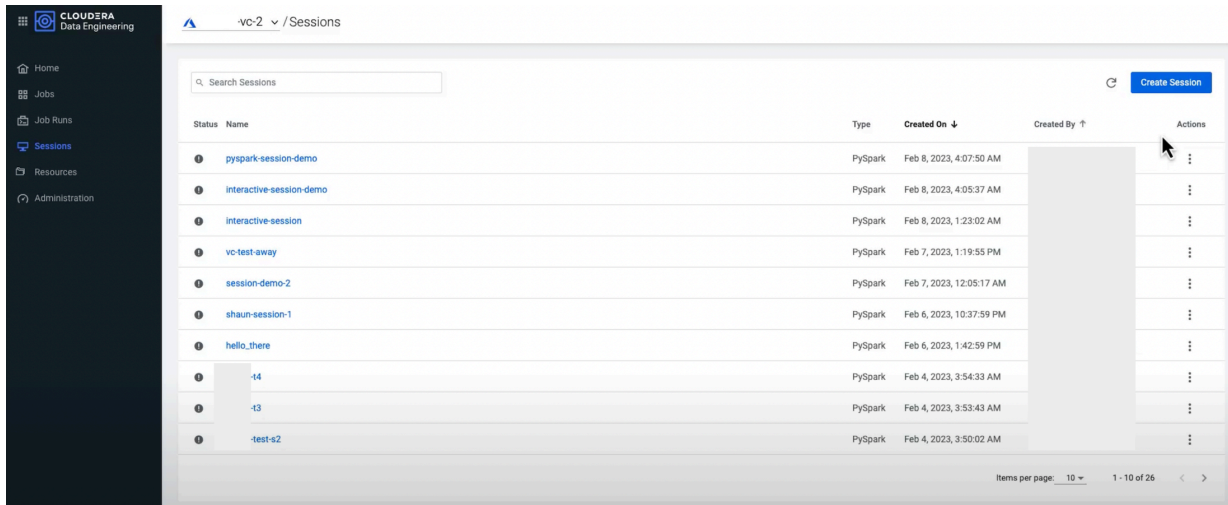
The commands that are run in a Cloudera Data Engineering Session are called Statements. You can submit the Statements through the connect CLI command or the Interact tab in the Cloudera Data Engineering UI for a Session. Python and Scala are the supported Session types. Learn how to use Cloudera Data Engineering Sessions using the user interface.

Before you begin

Ensure that you are using a version of Cloudera Data Engineering 1.19 or higher for your Virtual Cluster.

Procedure

1. In the Cloudera console, click the Data Engineering tile. The Home page displays.
2. Click Sessions in the left navigation menu and then click Create Session.



3. Enter a Name for the Session.
4. Select a Type, for example, PySpark, Scala, or Spark Connect.
5. Select a Timeout value.
The Session will stop after the indicated time has passed.
6. Optionally, enter a Description for the Session.
7. Optionally, enter the Configurations.
8. Set the Compute options.

9. Under **Files and Resources**, you can upload Jar, Python, Egg, Zip, and other files. You can also add a resource, repositories, or a Python environment to be used in this session.

Files that are uploaded to a session are stored in the app/mount directory.

Files and Resources
Add files and mount resources to be used in this session

Files & Resources

Jar Files

Upload or Select from Resource

Python, Egg, Zip files

Upload or Select from Resource

Other Files

Upload or Select from Resource

Resources ⓘ

Add Resource

Repositories ⓘ

Add Repository

Python Environment ⓘ

Select Python Environment

Cancel Create

10. Click Create.

The Connect tab displays a list of connectivity options available to interact with the Session. The Interact tab allows you to interact with the Session, and becomes available once the Session is running.

11. To delete a Session, open the Session and click Delete.



Note: If you delete a Session, doing so will result in the termination of an active session and the loss of any attached logs and details.

Interacting with a Session in Cloudera Data Engineering

Once a Session is created in Cloudera Data Engineering, you can begin your interaction.

About this task

You can also interact with files that were uploaded to a Session. Those files are stored in the app/mount directory.

Before you begin

Ensure that you are using a version of Cloudera Data Engineering 1.19 or higher for your Virtual Cluster.

Procedure

1. In the Cloudera console, click the Data Engineering tile. The Home page displays.
2. Click Sessions in the left navigation menu.
3. Click the Interact tab.
4. In the Input box, enter your commands.
5. Click Run. A history of executed commands displays in the sequence of their execution.

Connecting Sessions with the CDE CLI

Once a Session is created in Cloudera Data Engineering, you can download the CDE CLI from the user interface then use the terminal on your local machine to set up the interactive session using the CDE CLI.

Before you begin

Ensure that you are using a version of Cloudera Data Engineering 1.19 or higher for your Virtual Cluster.

Procedure

1. In the Cloudera console, click the Data Engineering tile. The Home page displays.
2. Click Sessions in the left navigation menu.
3. Click the Connect tab.
4. Download the CDE CLI.
5. Use the CDE CLI to run and connect interactive commands:
 - a) Copy the command that displays.
 - b) Paste the command in the CDE CLI and continue to interact with your session.

External IDE connectivity through Spark Connect-based sessions (Technical Preview)

You can learn what an external IDE Spark Connect session is, certain known limitations, and the supported Runtime component versions.

What an external IDE Spark Connect session is

A session is an interactive short-lived development environment for running Spark commands. A Spark Connect Session is a type of [Cloudera Data Engineering Session](#) that exposes the Spark Connect interface. A Spark Connect Session allows you to connect to Spark from any remote Python environment.

Spark Connect allows you to connect remotely to the Spark clusters. Spark Connect is an API that uses the DataFrame API and unresolved logical plans as the protocol. The separation between client and server allows Spark and its open ecosystem to be leveraged from everywhere. It can be embedded in modern data applications, in IDEs and Notebooks. For more information about Spark Connect, identify the Spark version in your Virtual Cluster, and navigate to the relevant *Spark Connect Overview* page linked to that Spark version in the [Spark documentation](#).

Supported versions of Cloudera Runtime components

Ensure that you are using Spark 3.5.1 before you use Spark Connect Sessions.

Limitations

Spark Connect Sessions do not support the following:

- [Profile support](#): Spark Connect does not support profiles in the configuration files even though the Cloudera Data Engineering clients support "Profiles" in the configuration files.
- **PySpark**: In Spark 3.4, Spark Connect supports most PySpark APIs, including DataFrame, Functions, and Column. Some APIs, such as SparkContext and RDD are not supported. You can check which APIs are currently supported in the [Apache Spark API Reference](#) documentation. Supported APIs are labeled "Supports Spark Connect", so before migrating existing code to Spark Connect, you can check whether the APIs you are using are available. For more information, see the [Apache Spark documentation](#).
- **Scala**: In Spark 3.5, Spark Connect supports most Scala APIs, including Dataset, functions, Column, Catalog, and KeyValueGroupedDataset. For more information, see the [Apache Spark documentation](#).
- User-Defined Functions (UDFs) are supported, by default, for the shell and in standalone applications, with additional setup requirements.
- The majority of the Streaming API is supported, including DataStreamReader, DataStreamWriter, StreamingQuery, and StreamingQueryListener. For more information, see the [Apache Spark documentation](#).
- APIs, such as SparkContext and RDD are deprecated in all Spark Connect versions.

Configuring external IDE Spark Connect sessions

Learn about how to configure a Spark Connect Session with Cloudera Data Engineering.

Before you begin

Before you create a Spark Connect Session, perform the following steps:

1. [Enable a Cloudera Data Engineering service](#) .
2. [Create a Cloudera Data Engineering Virtual cluster](#). You must select All Purpose (Tier 2) in the Virtual Cluster option and Spark 3.5.1 as the Spark version.

Procedure

1. Perform the following steps on each user's machine:

- a) Create the `~/.cde/config.yaml` configuration file and add the `vcluster-endpoint` and `cdp-endpoint` parameters.

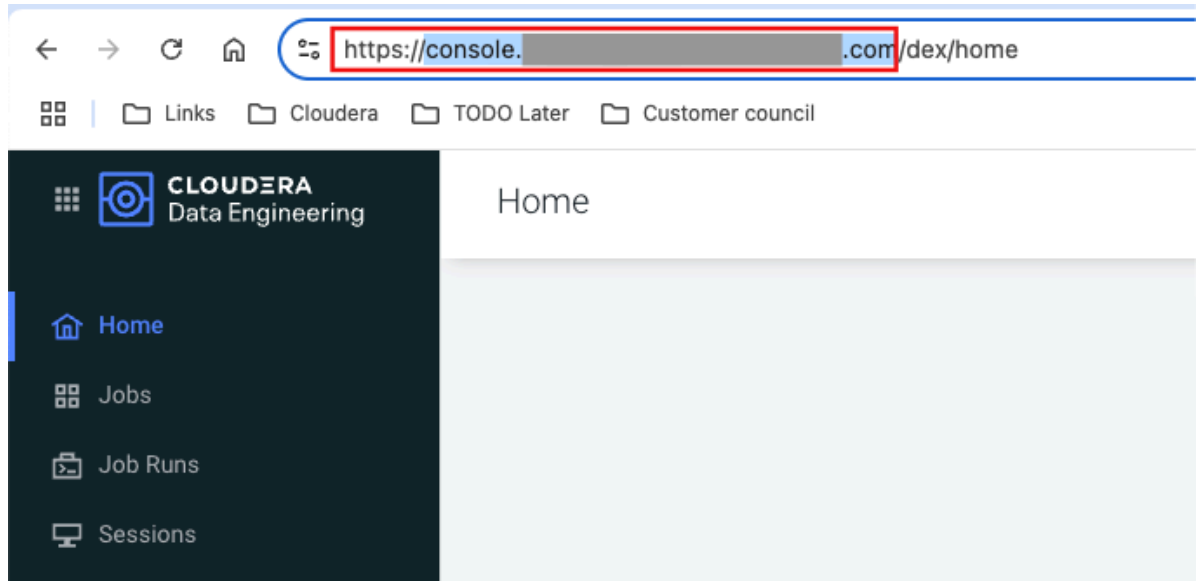
This allows the client machine to identify a virtual cluster.



Note: The `cdp-endpoint` value is the same as the Cloudera console URL. From the Cloudera console URL, copy the protocol (`https://`), the subdomain (`console`), the second-level domain, and the top-level domain (`.com`).

Example: `https://console.[***SECOND-LEVEL-DOMAIN***].com`

Figure 1: Getting the Cloudera endpoint URL from the Cloudera console URL



For more information, see [vcluster-endpoint](#) and [cdp-endpoint](#).

For example,

```
cdp-endpoint: https://console.cdp.apps.example.com
credentials-file: /Users/user1/.cde/credentials
vcluster-endpoint: https://ffws6v27.cde-c9b822vr.apps.example.com/dex/
api/v1
```

- b) [Create an access key](#) and update the `credentials-file` parameter in the `~/.cde/config.yaml` configuration file with the path where the credentials file is located. This allows the client machine to acquire the short-lived access tokens.



Note: Access keys configured with the default profile are supported.

For example,

```
[default]
cdp_access_key_id=571ff....
cdp_private_key=dvbYd....
```


2. Create a Spark Connect Session using one of the following methods:



Note: You can interact with a Spark Connect session that only you have created.

- Using the UI: Create a new session as per [Creating Sessions in Cloudera Data Engineering](#) but when you select the session type, select Spark Connect (Tech Preview) from the Type drop-down list.

Name *

Session Name

Type *

Spark Connect (Tech Preview)

Description

Describe the session

Configurations (Optional)

config_key config_value

Compute Options

Configure the compute options for this session

Number Executors	<input type="range"/>	1 40	<input type="text" value="1"/>
Driver Cores	<input type="range"/>	1 16	<input type="text" value="1"/>
Executor Cores	<input type="range"/>	1 16	<input type="text" value="1"/>
Driver Memory (GB)	<input type="range"/>	1 32	<input type="text" value="1"/>
Executor Memory (GB)	<input type="range"/>	1 32	<input type="text" value="1"/>

Cancel Create

- [Using the CLI](#): Create a Spark Connect Session by running the following command:

```
cde session create --name [***SPARK-SESSION-NAME***] --type spark-connect
```



Note:

To get all the attributes of a cde session command, run the `cde session -h` command.

3. On the Cloudera Data Engineering Home page, click Sessions and then select the Spark Connect Session that you have created.
4. Go to the Connect tab and download the required Cloudera Data Engineering TAR file and PySpark TAR file as displayed on the screen.

The screenshot shows the Cloudera Data Engineering (CDE) interface for a session named 'test-connect-2'. At the top, there's a table with session details: STATUS (Available), TYPE (Spark Connect), CREATED BY (cdpuser1), START TIME (Jan 31, 2024, 2:34:35 PM), and TIMEOUT (-). Below this is a tabbed interface with 'Configuration', 'Connect', and 'Logs'. The 'Connect' tab is active, displaying 'Connect with Spark Connect' instructions. It includes a list of steps: 1. Download the required CDE Tarball file to work with your Spark Connect Session; 2. Download the Pyspark 3.4 tarball required by the CDESparkConnectSession package; 3. Install in your python environment by running a code block:

```
pip install <cde connect tarball>
pip install <pyspark tarball>
```

; 4. Configure the python package. Below this is 'Step 2 : Connect' with step 1: Use the package to connect with this session, followed by a code block:

```
from cde import CDESparkConnectSession
spark = CDESparkConnectSession.builder.sessionName('test-connect-2').get()
```

. At the bottom, it says 'Need help? Follow our detailed guide' with a link icon.



Note:

- The Copy Link option can be used to retrieve a URL and download the client using cURL.
- The PySpark TAR file version must be same as the Virtual Cluster's Spark version.

5. Create a new Python virtual environment or use your existing one and install the TAR file after activating your Python virtual environment.

```
python3 -m venv cdeconnect
. cdeconnect/bin/activate

pip install [***CDECONNECT_TARBALL***]
pip install [***PYSPARK_TARBALL***]
```

Sample code to connect to an external IDE Spark Connect session

After configuring Spark Connect Sessions, learn how you can run the CLI commands from a remote Python host to connect to a session and execute Spark SQL commands through an example.

You can use the following sample code to connect to the Spark Connect session. Use the spark variable to interact with Spark as you connect to the Cloudera Data Engineering jobs or sessions.

```
> python
Python 3.9.13 (main, Jul 29 2022, 12:22:24)
[Clang 13.0.0 (clang-1300.0.27.3)] on darwin
Type "help", "copyright", "credits" or "license" for more information.
>>> from cde import CDESparkConnectSession
>>> spark = CDESparkConnectSession.builder.sessionName('connect-session').ge
t()
>>> spark.version
'3.4.1.1.20.7180.0-33'
>>> spark.sql("use retaildb").show()
++
||
++
++

>>> spark.sql("select * from products_external").show()
+-----+-----+-----+-----+
|product_id|product_category_id|product_name|product_description|pro
duct_price|product_image|
+-----+-----+-----+-----+
|1|2|Quest Q64 10 FT. ...|
59.98|http://images.acm...|
|2|2|Under Armour Men'...|
129.99|http://images.acm...|
|3|2|Under Armour Men'...|
89.99|http://images.acm...|
|4|2|Under Armour Men'...|
89.99|http://images.acm...|
|5|2|Riddell Youth Rev...|
199.99|http://images.acm...|
|6|2|Jordan Men's VI R...|
134.99|http://images.acm...|
|7|2|Schutt Youth Recr...|
99.99|http://images.acm...|
|8|2|Nike Men's Vapor ...|
129.99|http://images.acm...|
|9|2|Nike Adult Vapor ...|
50.0|http://images.acm...|
|10|2|Under Armour Men'...|
129.99|http://images.acm...|
|11|2|Fitness Gear 300 ...|
209.99|http://images.acm...|
|12|2|Under Armour Men'...|
139.99|http://images.acm...|
|13|2|Under Armour Men'...|
89.99|http://images.acm...|
|14|2|Quik Shade Summit...|
199.99|http://images.acm...|
|15|2|Under Armour Kids...|
59.99|http://images.acm...|
|16|2|Riddell Youth 360...|
299.99|http://images.acm...|
|17|2|Under Armour Men'...|
129.99|http://images.acm...|
|18|2|Reebok Men's Full...|
29.97|http://images.acm...|
|19|2|Nike Men's Finger...|
124.99|http://images.acm...
```

```
|      20 |      2 | Under Armour Men' ... |
| 129.99 | http://images.acm... |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
only showing top 20 rows
```

Troubleshooting errors when working with an external IDE Spark Connect session

While working with the Spark Connect Sessions in Cloudera Data Engineering, you might encounter errors. Learn how you can troubleshoot those errors.

Condition

If the session is killed or the driver exits due to an error when the code is being executed, Spark Connect shows the following error.

```
pyspark.errors.exceptions.connect.SparkConnectGrpcException: <_MultiThreaded
Rendezvous of RPC that terminated with:
  status = StatusCode.UNKNOWN
  details = "Stream removed"
  debug_error_string = "UNKNOWN:Error received from peer {grpc_message:"Str
eam removed", grpc_status:2, created_time:"2024-01-31T13:28:23.35214+05:30"}
"
>
```

Remedy

Procedure

Check the actual error from the session driver logs using [UI](#) or [CDE CLI](#).

Viewing logs for Cloudera Data Engineering Sessions

Once a Session is created in Cloudera Data Engineering, you can view logs of the Spark driver and executors from the Session.

Before you begin

Ensure that you are using a version of Cloudera Data Engineering 1.19 or higher for your Virtual Cluster.

Procedure

1. In the Cloudera console, click the Data Engineering tile. The Home page displays.
2. Click Sessions in the left navigation menu.
3. Click the Session name that you want to view the logs for.
4. Click the Logs tab.
5. In the Select Log Type drop-down, select the Driver or Executor that you want to view the logs for.
6. Click Download All Logs to save your log files.

Viewing Spark UI for Cloudera Data Engineering Sessions

You can view the Spark UI for a Cloudera Data Engineering active Session for troubleshooting purposes.

About this task



Important: Spark UI is not available for terminated sessions.

Procedure

1. In the Cloudera console, click the Data Engineering tile. The Home page displays.
2. Click Sessions in the left navigation menu.
3. Select the Virtual Cluster from the drop-down menu at the top.
4. Select the Active/Running sessions that you want to troubleshoot.
5. Click the Spark UI tab.