

Using Cloudera Data Engineering resources

Date published: 2020-07-30

Date modified: 2024-03-05

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with a stylized 'E' that has three horizontal bars.

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

Using Python virtual environments with Cloudera Data Engineering.....	4
Creating a Python virtual environment resource.....	4
Using credentials for custom pip repositories.....	6
Associating a Python virtual environment with a Cloudera Data Engineering job.....	7
Updating Python virtual environment resources.....	8
 Using Custom Spark Runtime Docker Images Via API/CLI.....	 9

Using Python virtual environments with Cloudera Data Engineering

Cloudera Data Engineering (CDE) supports Python virtual environments to manage job dependencies by using the `python-env resource` type.

A *resource* in CDE is a named collection of files or other resources referenced by a job. The `python-env resource` type allows you to specify a `requirements.txt` file that defines a virtual environment that you can then associate with a CDE job. You can specify any Python package in the `requirements.txt` file, including those with C dependencies.

Creating a Python virtual environment resource

After you have created the `requirements.txt` file, you can create the Python virtual environment resource.



Note:

- For `python-env` resources, you can only upload a `requirements.txt` file. Python environment resources do not support arbitrary file upload. If the local file is named something other than `requirements.txt`, you must add the flag `--resource-path requirements.txt` to the command, which renames the file to `requirements.txt` in the resource.

You can also specify a PyPi mirror for a Python virtual environment resource using the `--pypi-mirror` flag. This requires network access to the mirror from the CDP environment.

- If a Python package specified in the `requirements.txt` file is not found or does not support the underlying OS architecture `amd64`, then the Python virtual environment build fails.

For CDE CLI

Before you begin

- Download and configure the CDE CLI.
- Create a `requirements.txt` file specifying the Python package and version dependencies required by your CDE job.
- Ensure that the following hostnames are reachable from within the cluster, to install the Python package successfully if no PyPi mirror is configured:
 - `pypi.python.org`
 - `pypi.org`
 - `pythonhosted.org`
 - `files.pythonhosted.org`

Steps

1. Run the `cde resource create` command as follows to create a Python virtual environment resource.

```
cde resource create --name <***RESOURCE_NAME***> --type <***ENVIRONMENT_RESOURCE_TYPE***>
```

For example:

```
cde resource create --name cde-python-env-resource --type python-env
```



Note:

- You can use specific credentials for your custom pip repositories. For more information, see [Using credentials for custom pip repositories](#).
- You can use the optional flags when creating an Python environment. For more information, see [CDE CLI custom Airflow Python environment flag reference](#).
- [Optional] You can specify the custom pip repository using the `--pip-repository-url <***CUSTOM-PIP-REPOSITORY-URL***> --pip-repository-cert <***PATH-TO-PEM-FILE***>` option in the create resource command.
- [Optional] You can specify one or more extra custom pip repositories using the `--extra-pip-repository-<***NUMBER***>-url --<***CUSTOM-PIP-REPOSITORY-URL***>-<***NUMBER***>-cert <***PATH-TO-PEM-FILE***>` option in the create resource command. You can specify up to 10 extra pip repositories.

Example of command with pip repository and extra pip repository:

```
cde resource create --name cde-python-env-resource --type python-env --pip-repository-url https://pypi.example.com/simple --pip-repository-cert cert.pem --extra-pip-repository-1-url https://extra-pypi.example.com/simple --extra-pip-repository-1-cert extra-cert.pem
```

2. Upload the `requirements.txt` file to the resource.



Note: For `python-env` resources, you can only upload a `requirements.txt` file. Python environment resources do not support arbitrary file upload. If the local file is named something other than `requirements.txt`, you must add the flag `--resource-path requirements.txt` to the command, which renames the file to `requirements.txt` in the resource.

```
cde resource upload --name cde-python-env-resource --local-path ${HOME}/requirements.txt
```

Result

When you first create a Python virtual environment resource, CDE builds the environment according to the `requirements.txt` file. During this build time, you cannot run a job associated with the virtual environment. You can check the status of the environment by running `cde resource list-events --name <resource_name>`. For example:

```
cde resource list-events --name cde-python-env-resource
```

The environment is ready when you see a message similar to the following:

```
{
  "id": 4,
  "message": "Job pp-84kgd6f6-resource-builder-cde-python-env-resource-1634911572 succeeded, marking resource with ready status",
  "created": "2021-10-22T14:09:13Z"
}
```

For Web UI

Before you begin

- Create a requirements.txt file specifying the Python package and version dependencies required by your CDE job.
- Ensure that the following hostnames are reachable from within the cluster, to install the Python package successfully if no PyPi mirror is configured:
 - pypi.python.org
 - pypi.org
 - pythonhosted.org
 - files.pythonhosted.org

Steps

1. In the Cloudera Data Platform (CDP) console, click the Data Engineering tile. The CDE Home page displays.
2. Click Resources in the left navigation menu and then click Create Resource.
3. Specify a resource name, and then select Python Environment from the Type drop-down menu.
4. Choose the Python version for the environment and optionally specify the PyPi Mirror URL. The PyPi mirror must be accessible from the CDP environment.
5. Click Create.
6. Click Upload File and select the requirements.txt file from your local machine. You can also drag-and-drop the file to the outlined area on the page.

Result

The UI displays Building the resource... while the Python virtual environment is building. After the environment is built, the page displays the Python packages and versions included in the environment.

Using credentials for custom pip repositories

You can use specific credentials for your custom pip repositories. You can create a credential of type `basic` and use it when creating a Python environment resource. You must use this credential name to authenticate against the pip repository for the `python-env` or `airflow-python-env` resource.

```
cde credential create --type basic --name <user-credential> --username <user-name>
```

When prompted, provide the password for the specified user name.

Example to create and use the credential with the pip repository.

- For `python-env` resource:

```
cde credential create --type basic --name testuser-cred --username userdemo
cde resource create --name test-pyenv-with-cred --type python-env --pip-repository-url https://pypi.example.com/simple --pip-repository-cred testuser-cred --pip-repository-cert ~/Downloads/cert.pem
```

- For `airflow-python-env` resource:

```
cde credential create --type basic --name testuser-cred --username userdemo
cde resource create --name test-airflow-pyenv-with-cred --type airflow-python-env --pip-repository-url https://pypi.example.com/simple --pip-repository-cred testuser-cred --pip-repository-cert ~/Downloads/cert.pem
```

Associating a Python virtual environment with a Cloudera Data Engineering job

You can associate the Python virtual environment with a CDE job at the time of creation, or you can update an existing job.

For CDE CLI

Before you begin

- Download and configure the CDE CLI.
- Create a Python virtual environment CDE resource.
- Create a CDE job.

Steps

1. Using the CDE CLI, run the `cde job update` command to associate a Python virtual environment with the job.

```
cde job update --name pyspark-example --python-env-resource-name cde-python-env-resource
```



Note: You can specify a Python virtual environment resource at job creation time as well, using the flag `--python-env-resource-name`. For example:

```
cde job create --type spark --application-file pyspark-example.py --python-env-resource-name cde-python-env-resource --name pyspark-example
```

For Web UI

Before you begin

- Create a Python virtual environment CDE resource.
- Create a CDE job.

Steps

1. In the Cloudera Data Platform (CDP) console, click the Data Engineering tile. The CDE Home page displays.
2. Click **Job Runs** on the left navigation menu. The Job Runs page displays.

- Using the dropdown menu, select the virtual cluster containing the application you want to manage.

Status	Run ID
✓	1
✓	2
!	3

- Click on the name of the job you want to modify.
- Go to the Configuration tab.
- Click Edit.
- In the Python Environment section, click Select Python Environment.
- Select the Python virtual environment resource you want to use, and then click Select Resource.
- At the bottom of the page, click Update and Run to run the job immediately, or click the drop-down arrow on the button and select Update to update the job without running it.

Updating Python virtual environment resources

Currently, Python virtual environments cannot be updated. Instead, create a new Python virtual environment resource and update the job to reference the new resource.

For CDE CLI

Before you begin

- Download and configure the CDE CLI.
- Create a new requirements.txt file specifying the Python package and version dependencies required by your CDE job.

Steps

1. Create a new python-env resource.

```
cde resource create --name new-cde-python-env-resource --type python-env
--python-version python3
```

2. Upload the new or updated requirements.txt file to the new resource.

```
cde resource upload --name new-cde-python-env-resource --local-path ${HOME}/requirements.txt
```

3. Update the CDE job to specify the new resource.

```
cde job update --name pyspark-example --python-env-resource-name new-cde-python-env-resource
```

For Web UI

Before you begin

- Create a new requirements.txt file specifying the Python package and version dependencies required by your CDE job.

Steps

1. Create a new Python virtual environment using the new requirements.txt file, following the instructions in [Creating a Python virtual environment resource](#).
2. Update the CDE job to reference the new Python virtual environment, following the instructions in [Associating a Python virtual environment with a CDE job](#).

Related Information

[Creating a Python virtual environment resource](#)

[Associating a Python virtual environment with a Cloudera Data Engineering job](#)

Using Custom Spark Runtime Docker Images Via API/CLI

This is a detailed usage guide to demonstrate how to run jobs using custom spark runtime with examples.

Steps

1. Create a custom docker image

Build “custom-spark-dex-runtime” images based on the dex-spark-runtime image of the CDE version.



Note: The image should be based on the dex-spark-runtime of the current dex version.

The relevant dex-spark-runtime image is

```
<registry-host>/cloudera/dex/dex-spark-runtime-<spark version>-<cdh version>:<dex version>
```

Example: DockerFile for DEX 1.18.2-b70, Spark 2.4.7 version, and CDH version.

```
DockerFile for DEX 1.18.2-b70, spark 2.4.7 version, and CDH version.
```

```
FROM docker-private.infra.cloudera.com/cloudera/dex/dex-spark-runtime-2
.4.7-7.1.7.1000:1.18.2-b70
USER root
RUN yum install -y git && yum clean all && rm -rf /var/cache/yum
RUN pip2 install virtualenv-api
RUN pip3 install virtualenv-api
USER ${DEX_UID}
```

2. Build the docker image tagging it with the custom registry to be used and push it to the custom registry.

Example:

```
mac@local:$ docker build --network=host -t docker.my-company.registry.com/
custom-dex/dex-spark-runtime-2.4.7-7.1.7.1000:1.18.2-b70-custom . -f Dockerfile
mac@local:$ docker push docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.7-7.1.7.1000:1.18.2-b70-custom
```

Here, the custom registry is `docker.my-company.registry.com` and the registry namespace is `custom-dex`.



Note: Obtain `$CDE_TOKEN` to execute the REST API examples by following [Getting a Cloudera Data Engineering API access token](#) document.

3. Create the credentials for the custom image registry.

Register docker registry image pull credentials using the CDE CLI or REST API. These credentials are stored as a secret.



Note: If using public registry skip to this [step](#)

For CLI

```
mac@local:$ ./cde credential create --name docker-creds --type docker-basic --docker-server docker-sandbox.infra.cloudera.com --docker-username my-username
```

For REST API

```
curl -X POST -k 'https://<dex-vc-host>/dex/api/v1/credentials' \
-H "Authorization: Bearer ${CDE_TOKEN}" \
-H 'accept: application/json' \
-H 'Content-Type: application/json' \
--data '{
  "dockerBasic": {
    "password": "password123",
    "server": "docker-sandbox.infra.cloudera.com",
    "username": "my-username"
  },
  "name": "docker-creds",
  "type": "docker-basic"
}'
```

4. Create custom-runtime-image resource referring to the credential created previously.

Register “custom-spark-dex-runtime” docker image as a resource of type custom-runtime-image specifying the name of the credential created in the previous step.

For CLI

```
mac@local:$ ./cde resource create --name custom-image-resource --image
```

```
docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.7-7.1.7.1000:1.18.2-b70-custom --image-engine spark2 --type custom-runtime-image --image-credential docker-creds
```

For REST API

```
curl -X POST -k 'https://<dex-vc-host>/dex/api/v1/resources \
-H "Authorization: Bearer ${CDE_TOKEN}" \
-H 'accept: application/json' \
-H 'Content-Type: application/json' \
--data '{
  "customRuntimeImage": {
    "credential": "docker-creds",
    "engine": "spark2",
    "image": "docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.7-7.1.7.1000:1.18.2-b70-custom"
  },
  "name": "custom-image-resource",
  "type": "custom-runtime-image"
}'
```

5. Submit a job by setting the “custom-spark-dex-runtime” image as a resource using the CLI

For SPARK COMMAND

```
mac@local:$ ./cde --user cdpuser1 spark submit /Users/my-username/spark-examples_2.11-2.4.4.jar
--class org.apache.spark.examples.SparkPi 1000 --runtime-image-resource-name=custom-image-resource
```

For JOB COMMAND

```
mac@local:$ ./cde --user cdpuser1 resource create --name spark-jar
mac@local:$ ./cde --user cdpuser1 resource upload --name spark-jar --local-path spark-examples_2.11-2.4.4.jar

mac@local:$ ./cde --user cdpuser1 job create --name spark-pi-job-cli --type spark --mount-1-resource spark-jar --application-file spark-examples_2.11-2.4.4.jar --class org.apache.spark.examples.SparkPi --user cdpuser1 --arg 22 --runtime-image-resource-name custom-image-resource
```

6. The spark driver/executor pods should use this specific image and you can confirm it by opening a shell into those pods and verifying if the external installed libraries or files exist.

Public docker registries

Create the resource for the registries which do not require any auth. You do not need to specify the credentials.

For CLI

```
mac@local:$ cde resource create --name custom-image-resource --image docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.7-7.1.7.1000:1.18.2-b70-custom --image-engine spark2 --type custom-runtime-image
```

For REST API

```
ccurl -X POST -k 'https://<dex-vc-host>/dex/api/v1/resources \
-H "Authorization: Bearer ${CDE_TOKEN}" \
```

```
-H 'accept: application/json' \  
-H 'Content-Type: application/json' \  
--data '{  
  "customRuntimeImage": {  
    "engine": "spark2",  
    "image": "docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.7-7.1.7.1000:1.18.2-b70-custom"  
  },  
  "name": "custom-image-resource",  
  "type": "custom-runtime-image"  
}'
```

Once done, skip to #step 5 to submit the job.

Error: Custom image resource with missing or wrong credentials

Creating a custom image resource with missing or wrong credentials should result in the below error which can be seen in the logs or in kubernetes pod events.

Example

```
Failed to pull image "docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.7-7.1.7.1000:1.18.2-b70-custom":  
rpc error: code = Unknown desc = Error reading manifest 1.18.2-b70-custom in  
docker.my-company.registry.com/custom-dex/dex-spark-runtime-2.4.7-7.1.7.1000:  
errors: denied: requested access to the resource is denied unauthorized: authentication required
```