Cloudera Data Engineering 1.25.0

Accessing the Cloudera Data Engineering service using the CLI

Date published: 2020-07-30
Date modified: 2025-11-20

CLOUDERA

Legal Notice

Contents

Using the Cloudera Data Engineering command line interface

Cloudera Data Engineering provides a command line interface (CLI) client. You can use the CLI to create and update jobs, view job details, manage job resources, run jobs, and so on.

> **Note:** The CLI client is not forward compatible. Download the client for the version of the cluster you are accessing. The Cluster Details page for every virtual cluster includes a link to download the CLI client for that cluster version.

The CLI client can also use a password file for non-interactive uses, such as automation frameworks.

Related Information

Using CLI-API to Automate Access to Cloudera Data Engineering

Using Cloudera Data Engineering CLI

Downloading the Cloudera Data Engineering command line interface

Cloudera Data Engineering provides a command line interface (CLI) client.

In addition to the CDE API, you can use the CDE CLI client to access your Cloudera Data Engineering service. Using the CLI, you can manage clusters and applications.

> **Note:** The CLI client is not forward compatible. Download the client for the version of the cluster you are accessing. The Cluster Details page for every virtual cluster includes a link to download the CLI client for that cluster version.

To download the CLI client:

1. Navigate to the Cloudera Data Engineering Overview page by clicking the Data Engineering tile in the Cloudera management console.
2. In the Cloudera Data Engineering web console, select an environment.
3. Click the Cluster Details icon for the virtual cluster you want to access.
4. Click the link under CLI TOOL to download the CLI client.
5. Change the permissions on the downloaded cde file to make it executable:

```
chmod +x /PATH/TO/cde
```

Configuring the CLI client

The CDE CLI client uses a configuration file, ~/.cde/config.yaml, to define the default Cloudera Data Engineering virtual cluster to interact with, as well as other configuration parameters.

Before you begin

Make sure that you have downloaded the CDE CLI client.

Procedure

1. Determine the virtual cluster endpoint URL.
    a) Navigate to the Cloudera Data Engineering Overview page.
    b) In the Environments column, select the environment containing the virtual cluster you want to access using the CLI.
    c) In the Virtual Clusters column on the right, click the Cluster Details icon for the virtual cluster you want to access.
    d) Click JOBS API URL to copy the URL to your clipboard.

2. On the host with the CLI client, create or edit the configuration file at ~/.cde/config.yaml. You can create multiple profiles in the ~/.cde/config.yaml file and can be used while running commands. For more information, see Creating and using multiple profiles.

    > **Note:** You can use a custom file location by setting the CDE_CONFIG environment variable. If you include "user: " in your config.yaml, the CLI tool will not use the access token and will instead keep prompting for a password. If you want to use a credential file, then you must remove the specified user from the profile.

3. In the configuration file, specify the Cloudera user and virtual cluster endpoint as follows:

    ```
    user: <CDP_USER>
    vcluster-endpoint: <CDE_VIRTUAL_CLUSTER_ENDPOINT>
    ```

    The Cloudera user is your workload username.

4. Save the configuration file.

5. If you have not done so already, make sure that the cde file is executable by running chmod +x   /PATH/TO/cde.

6. Run cde job list to verify your configuration. Enter your workload password when prompted.

    > **Note:** If the directory containing the cde file is not part of your PATH environment variable, you can either add it to your PATH environment variable or use the full path when running the command.

    > You can also configure the CLI to use an access token so that you do not need to enter your password each time. For more information, see Cloudera Data Engineering CLI authentication .

What to do next
See for other configuration options.

Cloudera Data Engineering CLI configuration options

The Cloudera Data Engineering CLI can be configured using a configuration file, environment variables, or by command flags.

Configuration Option	Configuration File (~/.cde/config.yaml)	Environment Variable	Command Flag
User	user: *<USERNAME>*	CDE_USER=*<USER>*	--user *<USERNAME>*
Credentials file	credentials-file: *</path/to/credentials>*	CDE_CREDENTIALS_FILE=*</path/to/credentials>*	--credentials-file *</path/to/credentials>*
Skip credentials file detection	skip-credentials-file: true	CDE_SKIP_CREDENTIALS_FILE=true	--skip-credentials-file
Password file	auth-pass-file: *<PASSWORD_FILE>*	CDE_AUTH_PASS_FILE=*<PASSWORD_FILE>*	--auth-pass-file *<PASSWORD_FILE>*
Virtual cluster endpoint	vcluster-endpoint: *<VIRTUAL_CLUSTER>*	CDE_VCLUSTER_ENDPOINT=*<VIRTUAL_CLUSTER>*	--vcluster-endpoint *<VIRTUAL_CLUSTER>*
Disable authentication token caching	auth-no-cache: true	CDE_AUTH_NO_CACHE=true	--auth-no-cache

Configuration Option	Configuration File (~/.cde/config.yaml)	Environment Variable	Command Flag
Authentication token cache file	auth-cache-file: *<TOKEN_CACHE_FILE>*	CDE_AUTH_CACHE_F ILE=*<TOKEN_CACHE_FILE>*	--auth-cach-file *<TOKEN_CACHE_FILE>*
CDE configuration profile		CDE_CONFIG_PROFI LE=*<PROFILE_NAME>*	--config-profile *<PROFILE_NAME>*

Creating and using multiple profiles using CDE CLI

You can now add a collection of CDE CLI configurations grouped together as profiles, to the config.yaml file. You can use these profiles while running commands. You can set the configurations either at a profile level or at a global level.

About this task

The CDE CLI client uses ~/.cde/config.yaml configuration file to define the default Cloudera Data Engineering virtual cluster to interact with and to define other configuration parameters. CDE CLI now supports the profiling option in the configuration file. A profile is a collection of configurations that you can apply to a CLI command. Multiple named profiles can be stored in the configuration file.

Before you begin

Make sure that you have downloaded the CLI client. For more information, see *Using the Cloudera Data Engineering command line interface*.

Procedure

1. Create or open the ~/.cde/config.yaml file.
2. Add profiles to the config.yaml file based on the following structure:

```
<Global Configurations>
profiles:
- name: <PROFILE NAME 1>
 <PROFILE CONFIGURATIONS>
- name: <PROFILE NAME 2>
 <PROFILE CONFIGURATIONS>
```

Example Configuration file:

```
vcluster-endpoint: https://g7f9bnv2.cde.dev.cldr.work/dex/api/v1
user: cdpuser

profiles:
- name: dev
  user: cdpuser1
  vcluster-endpoint: https://y86gbhv3.cde.dev.cldr.work/dex/api/v1

- name: test
  vcluster-endpoint: https://g7f9bnv2.cde8x.dev.cldr.work/dex/api/v1
```

- Global Configuration: These configurations are set at the global level. The configurations set here are used by default when a profile name is not specified or the configuration is not specified in the profile

- Profile Configuration: These configurations are set at the profile level and overrides the respective configurations set at the global level. You can select the profile either by using the environment variables, or command flags.

  > **Note:** Any profile mentioned under profiles overrides the keys of global configuration and does not replace all the configurations.
  >
  > For example, if there are five parameters in global and only three are configured in the selected profile, the final configurations will be three from the selected profile and the remaining from global . In the above test profile example, it does not have the user configuration defined, so if the user selects the test profile, except vcluster-endpoint, which is set in the test profile, all the other configurations are used from global .
  >
  > Global configuration is accessed as the default profile name. Hence, you cannot create a profile named default .

3. You can select the profile using the flag or environment variable.

   Flag

   ```
   cde job list --config-profile <PROFILE NAME>
   ```

   Environment variable

   ```
   export CDE_CONFIG_PROFILE=<PROFILE NAME>
   cde job list
   ```

4. [Optional] You can view the active profile in the configuration file.

   ```
   cde profile show-active
   ```

Cloudera Data Engineering CLI authentication

The Cloudera Data Engineering CLI tool supports both interactive and transparent authentication. For interactive authentication, if you have configured the CLI with your workload username, you are prompted for a password. For transparent authentication, the CDE CLI supports a password file, Cloudera access keys, and Cloudera credentials file .

The CDE CLI provides several mechanisms for authentication:

- Cloudera access key stored in a credentials file
- Cloudera access key specified by CLI flag or environment variable
- Interactive prompt for workload password
- Workload password specified by CLI flag or environment variable

> **Note:** If you include "user: " in your config.yaml, the CLI tool does not use the access token and instead keeps prompting for a password. If you want to use a credential file, then you must remove the specified user from the profile.

Caching

In all cases, the CLI uses the provided credentials to obtain an authentication token for the specified user, and caches it locally in a file on the machine where the CLI is running. You can disable caching of tokens entirely by using the --auth-no-cache CLI flag or the CDE_AUTH_NO_CACHE environment variable.

The cache file location is automatically determined based on the default system user cache:

- Linux: $HOME/.cache/cloudera/cde or $XDG_CACHE_HOME/cloudera/cde/
- macOS: $HOME/Library/Caches/cloudera/cde/
- Windows: %LocalAppData%\cloudera\cde\

If you want to use a custom location, specify it with the --auth-cache-file flag or the CDE_AUTH_CACHE_FILE environment variable. You can use the special string $USERCACHE, which is expanded according to the default system user cache (as listed above, without the /cloudera/cde/ suffix).

Cloudera credentials file

> ⚠️ **Important:** The minimum required roles to obtain an access token are DEUser and EnvironmentUser. EnvironmentAdmin role is not required.

When you generate a Cloudera access key, you can download it to a credentials file:
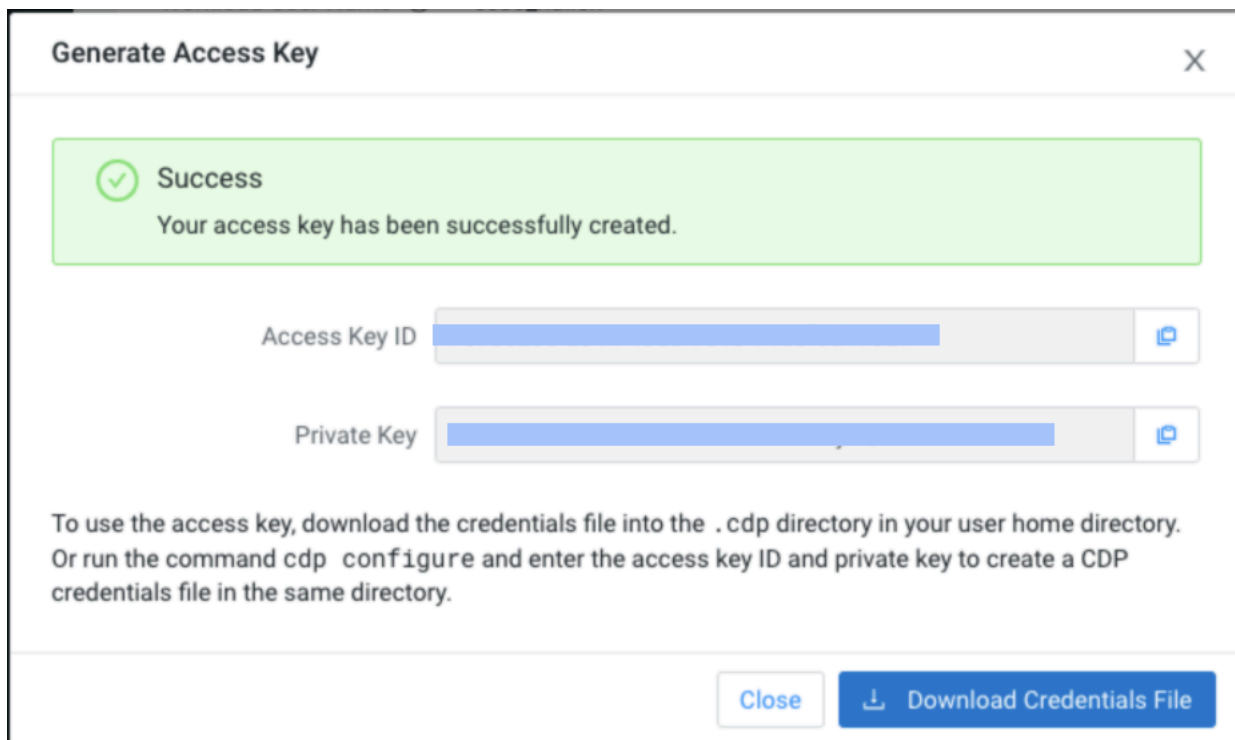
The access key is only displayed and available for download when you first generate it. After you close the dialog, there is no way to recover the key.

Save or copy the credentials file to $HOME/.cdp/credentials on the machine where you are running the CDE CLI. Credentials stored in this file are automatically discovered by both the CDE and CDP CLIs. If a credentials file is found, authentication occurs transparently using the discovered Cloudera access key.

The CDE CLI automatically looks for a Cloudera access key in the following locations in the order given:

1. ./credentials
2. $HOME/.cde/credentials
3. /etc/cde/credentials
4. $HOME/.cdp/credentials

> 📝 **Note:** If you include "user: " in your config.yaml, the CLI tool does not use the access token and instead keeps prompting for a password. If you want to use a credential file, then you must remove the specified user from the profile.

You can override this by using the --credentials-file *</PATH/TO/CREDENTIALS>* CLI flag to specify a different file location.

You can also skip credential discovery by using the --skip-credentials-file flag.

Cloudera access key

If you do not want to use the credentials file, you can specify the access key using environment variables or command line flags as follows:

Table 1: Cloudera access key environment variables and CLI flags

Parameter	Environment variable	CLI flag
Access key ID	CDE_ACCESS_KEY_ID=*<ACCESS_KEY_ID>*	--access-key-id *<ACCESS_KEY_ID>*
Access key secret	CDE_ACCESS_KEY_SECRET=*<ACCESS_KEY_SECRET>*	--access-key-secret string *<ACCESS_KEY_SECRET>*

Workload password prompt

When the CLI requires a new token for a virtual cluster, if you are not using a Cloudera access key, you are prompted for the password for the workload user, identified by the --user CLI flag or the CDE_USER environment variable.

The workload password, for both human and machine users, can be set using the Cloudera User Management console. For more information, see Cloudera workload user.

Workload password file

If you are not using a Cloudera access key, and you do not want to be prompted for your workload password, you can provide a password file. This is not the same as the credentials file discussed before. A password file is a file containing your workload password, and nothing else.

Note: When using a password file, the CLI strips one trailing newline character. If your password actually includes a newline character at the end, add an extra newline at the end of the file.

You can specify the password file by using an environment variable or a command line flag as follows:

Environment variable

> CDE_AUTH_PASS_FILE=*</PATH/TO/PASSWORD/FILE>*

Command line flag

> --auth-pass-file   *</PATH/TO/PASSWORD/FILE>*

Cloudera Data Engineering CLI TLS configuration

All Cloudera Data Engineering virtual cluster endpoints are configured with TLS. In non-production or on-premises environments the TLS certificates are usually signed by a non-production or non-public certificate authority (CA). In these cases, without additional configuration, the CLI tool fails as it attempts to validate the API server's TLS certificate. The CLI provides a TLS configuration when using non-public/non-production CAs.

Specify a file containing the PEM encoded public certificate(s) of the signing CA in one of the following ways:

- add the --tls-ca-certs *[***/PATH/TO/CA.PEM***]* flag on the command line
- define the tls-ca-certs: *[***/PATH/TO/CA.PEM***]* variable in the ~/.cde/config.yaml configuration file
- set the CDE_TLS_CA_CERTS environment variable

Replace *[***/PATH/TO/CA.PEM***]* with the path to a valid ca.pem file.

Certificates for Cloudera Data Engineering are issued and signed by LetsEncrypt and Amazon. If you are using an older system that does not already trust these certificate authorities, you must specify them as trusted using the --tls-ca-certs option. Additionally, if you are using the CDE CLI on Microsoft Windows with the --tls-ca-certs option, the operating system trusted root CAs are ignored. To make sure that your CA certificate file includes all of the necessary certificates, you can download the certificate chain from by doing the following:

1. Go to the Cloudera Data Engineering Overview page.
2. In the Environments column, select the environment containing the virtual cluster you want to interact with.
3. In the Virtual Clusters column on the right, click the Cluster Details icon for the virtual cluster you want to interact with.
4. Click JOBS API URL to copy the link to your clipboard.
5. Paste the link into your browser and press Enter.
6. In your browser, click the lock icon to view the certificate chain.
7. Export the certificate chain as a Base64-encoded X.509 certificate (.CER). Save it as ca.pem.

> **Tip:** When using the CLI on Windows, remember to use path styles such as C:\Users\janeblogs\.cde\ca.pem when referencing local files.

Cloudera Data Engineering concepts

Learn about some basic concepts behind Cloudera Data Engineering service to better understand how you can use the command line interface (CLI).

Cloudera Data Engineering has three main concepts:

job

A 'job' is a definition of something that Cloudera Data Engineering can run. For example, the information required to run a jar file on Spark with specific configurations.

job run

A 'job run' is an execution of a job. For example, one run of a Spark job on a Cloudera Data Engineering cluster.

session

A 'session' is an interactive short-lived development environment for running Spark commands to help you iterate upon and build your Spark workloads.

resource

A 'resource' refers to a job dependency that must be available to jobs at runtime. Currently the following resource types are supported:

- files is a directory of files that you can upload to Cloudera Data Engineering pods into a standard location (/app/mount). This is typically for application (for example, .jar or .py files) and reference files, and not the data that the job run will operate on. Multiple files resources can be referenced in a single job.
- python-env is used to provide custom Python dependencies to the job as a Python virtual environment which is automatically configured. Up to one python-env resource can be specified per job definition.
- custom-runtime-image is a custom Docker container image

> **Note:** Custom Docker container images is a *Technical Preview* feature. Contact your Cloudera account representative to enable access to this feature.

In addition, to support jobs with custom requirements, Cloudera Data Engineering also allows users to manage credentials which can be used at job run time. Currently, only custom Docker registry credentials are supported.

Submitting versus running a job

The cde spark submit and cde airflow submit commands automatically create a new job and a new resource, submit the job as a job run, and when the job run terminates they delete the job and resources.

A cde job run requires a job and all necessary resources to be created and uploaded to the Cloudera Data Engineering cluster beforehand. The advantage of creating resources and jobs ahead of time is that resources can be reused across jobs, and that jobs can be run using only a job name.

Managing Cloudera Data Engineering job resources using the CLI

A *resource* in Cloudera Data Engineering is a named collection of files or other resources referenced by a job. The files can include application code, configuration files, or any other arbitrary files required by a job. A resource can also be a Python virtual environment, or a custom Docker container image.

Note: Custom Docker container images is a *Technical Preview* feature. Contact your Cloudera account representative to enable access to this feature.

You can think of resources as any supporting files, libraries, or images that a Cloudera Data Engineering job requires to run. Resources can be created and deleted, and files can be added to and deleted from a resource as needed.

A resource can also be a Python virtual environment specification (as a requirements.txt file), or a custom Docker container image.

Before continuing, make sure that you have downloaded and configured the CLI client.

Creating a Cloudera Data Engineering resource using the CLI

A *resource* in Cloudera Data Engineering is a named collection of files or other assets referenced by a job, including application code, configuration files, or any other arbitrary files required by a job. A resource can also be a Python virtual environment, or a custom Docker container image.

Before you begin

Note: Custom Docker container images is a *Technical Preview* feature. Contact your Cloudera account representative to enable access to this feature.

Make sure that you have downloaded and configured the CLI client.

Procedure

1. Create a resource using the cde resource create command.

   The cde resource create syntax is as follows:

   ```
   cde resource create [flags]
   ```

   You can view the list of flags by running cde resource create   --help, or you can view the CDE CLI reference documentation.

   Example: Create a file resource

   ```
   cde resource create --name cde-file-resource --type files
   ```

   Example: Create a Python virtual environment resource

   ```
   cde resource create --name cde-python-env-resource --type python-env --p
   ython-version python3
   ```

   > **Note:**
   >
   > You can specify a PyPi mirror for a Python virtual environment resource using the --pypi-mirror flag. Note, that this requires network access to the mirror from the Cloudera environment.

   Example: Create a custom Docker container image resource

   ```
   cde resource create --name cde-container-image-resource --type custom-ru
   ntime-image
   ```

2. Verify that the resource was created by running cde   resource list.

Uploading files or other assets to a Cloudera Data Engineering resource using the CLI

A *resource* in Cloudera Data Engineering is a named collection of files or other assets referenced by a job, including application code, configuration files, or any other arbitrary files required by a job. A resource can also be a Python virtual environment, or a custom Docker container image.

Before you begin

> **Note:** Custom Docker container images is a *Technical Preview* feature. Contact your Cloudera account representative to enable access to this feature.

Make sure that you have downloaded and configured the CLI client.

Make sure that you have created a resource.

Procedure

1. Upload assets to a resource using the cde resource upload command.

   The cde resource upload syntax is as follows:

   ```
   cde resource upload [flags]
   ```

   You can view the list of flags by running cde resource upload --help, or you can view the CDE CLI reference documentation.

   > **Note:** For Python environment 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.

   Example: Upload a file resource

   ```
   cde resource upload --name [***RESOURCE_NAME***] --local-pat
   h [***LOCAL_PATH***] [--resource-path [***PATH_IN_RESOURCE***]]
   ```

   Use repeated local path flags, and/or */?/[a-z] wildcards, to specify multiple files. Use quotes around the local path when including wildcards, for example, --local-path "*.jar". For a single file --resource-path is used for the resource filename. For multiple files --resource-path is used for the resource directory.

   Example: Upload a Python virtual environment resource

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

   Example: Upload a custom Docker container image resource

   ```
   cde resource upload --name cde-container-image-resource --type custom-ru
   ntime-image
   ```

   Example: Upload a file for a file resource

   ```
   cde resource upload --name cde-file-resource --local-path /path/to/local/
   file
   ```

   Example: Upload and extract archive to resource

   Currently supported archive file formats are : .zip and .tar.gz

   ```
   cde resource upload-archive --name cde-file-resource --local-path /path/
   to/local/file
   ```

2. Verify that the file is included in the resource by running cde   resource describe --name *<RESOURCE_NAME>*.

Deleting a Cloudera Data Engineering resource using the CLI

A *resource* in Cloudera Data Engineering is a named collection of files or other resources referenced by a job, including application code, configuration files, or any other arbitrary files required by a job. A resource can also be a Python virtual environment, or a custom Docker container image. Resources can be deleted using the CLI.

Before you begin

- Make sure that you have downloaded and configured the CLI client.
- Make sure that the resource you are deleting is no longer needed for any jobs. (Resources cannot be deleted if they are referenced in one or more jobs)

Procedure

1. Run cde resource describe --name   *<RESOURCE_NAME>*. View the output and confirm that the resource you want to delete is no longer required, and does not contain any files that you need to retain.
2. Delete the resource by running cde resource delete   --name *<RESOURCE_NAME>*
3. Verify that the resource is deleted by running cde resource list and confirming that the resource is no longer listed.

Creating and updating credentials

To allow the use of private Docker registries, Cloudera Data Engineering supports the creation and management of credentials. Credentials are secure objects used to store access keys (such as passwords, keys, or tokens) that are required for various workflows within Cloudera Data Engineering. These workflows can involve accessing resources, and repositories. These credentials are stored securely in the Kubernetes cluster as secrets and cannot be accessed by end users directly. Credentials are attached to job runs automatically by the Cloudera Data Engineering backend.

About this task

> **Note:** Custom Docker container images is a *Technical Preview* feature. Contact your Cloudera account representative to enable access to this feature.

Procedure

1. To create a new credential:

```
cde credential create --name <CRED_NAME> --type docker-basic --docker-se
rver <REGISTRY_URL_OR_HOSTNAME> --docker-username <DOCKER_USER>
```

2. Enter the Docker registry password when you are prompted.

   An optional --description field allows you to annotate the credential with a human readable description.

3. Run cde credential list to verify that the credential was created:

```
cde credential list [--filter <FILTER>]
```

   For more information on filtering syntax, see CDE CLI list command syntax reference on page 37.

4. If you want to update a credential, use the cde credential update command.

   This command allows you to update the secret content, the credential description, or both.

```
cde credential update --name <CRED_NAME> [--docker-serve
r <REGISTRY_URL_OR_HOSTNAME> --docker-username <DOCKER_USER>] [--descrip
tion "<DESC>"]
```

Deleting credentials

To allow the use of private Docker registries, Cloudera Data Engineering supports the creation and management of credentials. Credentials are secure objects used to store access keys (such as passwords, keys, or tokens) that are required for various workflows within Cloudera Data Engineering. These workflows can involve accessing resources, and repositories. These credentials are stored securely in the Kubernetes cluster as secrets and cannot be accessed by end users directly. Credentials are attached to job runs automatically by the Cloudera Data Engineering backend.

Before you begin

- Make sure that you have downloaded and configured the CLI client.
- Make sure that the credential you are deleting is no longer needed for any jobs.

About this task

Note: Custom Docker container images is a *Technical Preview* feature. Contact your Cloudera account representative to enable access to this feature.

Procedure

1. Delete the credential by running cde credential delete --name   *<CRED_NAME>*
2. Run cde credential list to verify that the credential was deleted:

```
cde credential list [--filter <FILTER>]
```

For more information on filtering syntax, see CDE CLI list command syntax reference on page 37.

Deleting an Airflow DAG

You can delete unused Airflow DAGs using the Cloudera Data Engineering command line interface (CLI).

About this task
The default process of removing Cloudera Data Engineering resources is to delete them together with the jobs owning them, using the cde job delete command. The cde airflow delete-dag command is a fallback for when Airflow gets into an unexpected situation and you have to remove a DAG with no associated Airflow job.

Procedure

To delete a DAG from Airflow that is not associated with a job, use the cde airflow delete-dag command:

```
cde airflow delete-dag --dag-id <DAG_ID>
```

Managing Cloudera Data Engineering jobs using the CLI

A *job* in Cloudera Data Engineering is a definition of something that Cloudera Data Engineering can run. For example, the information required to run a JAR file on Spark with specific configurations. A 'job run' is an execution of a job. For example, one run of a Spark job on a Cloudera Data Engineering cluster.

Creating and updating Apache Spark jobs using the CLI

The following example demonstrates how to create a Spark application in Cloudera Data Engineering using the command line interface (CLI).

Before you begin
Make sure that you have downloaded the CLI client. For more information, see Using the Cloudera Data Engineering command line interface. While creating a job if you want to use the [--data-connector] flag, you must obtain the name of the data connector from the Cloudera Data Engineering UI by navigating to Administration > click Service Details icon of the Cloudera Data Engineering Service > Data Connectors tab.

Procedure

1. Run the cde job create command as follows:

```
cde job create --application-file <PATH_TO_APPLICATION_JAR> --mount-1-r
esource <YOUR_CDE_RESOURCE> --class <APPLICATION_CLASS> [--default-vari
```

```
able name=value] [--data-connector name] --name <JOB_NAME> --num-execu
tors <NUM_EXECUTORS> --type spark
```

To see the full command syntax and supported options, run cde job create   --help.

The application file can either be a local JAR file inside the Spark container or can be mounted from a Cloudera Data Engineering Resource. If it is from a Cloudera Data Engineering Resource, the --application-file points to the file from within the Cloudera Data Engineering Resource. If you want to configure the job to pick up the application file from a Cloudera Data Engineering Repository, provide the name of the repository with the --mo unt-1-resource flag and provide the --application-file path as you would do with a Cloudera Data Engineering Resource.

With [--default-variable] flags you can replace strings in job values. Currently the supported fields are:

- Spark application name
- Spark arguments
- Spark configurations

For a variable flag name=value any substring {{{name}}} in the value of the supported field gets replaced with value. These can be overriden by the [--variable] flag during the job run.

Using the [--data-connector] flag, you can specify the name of the data connector. Currently, only the Ozone type data connector is supported and it must be created before the job run.

2. Run cde job describe to verify that the job was created:

```
cde job describe --name <JOB_NAME>
```

3. If you want to update the job configuration, use the cde job update command.
   For example, to change the number of executors:

```
cde job update --name test_job --num-executors 15
```

To see the full command syntax and supported options, run cde job update   --help.

4. To verify the updated configuration, run cde job   describe again:

```
cde job describe --name <JOB_NAME>
```

Creating and updating Apache Airflow jobs using the CLI

The following example demonstrates how to create an Airflow DAG in Cloudera Data Engineering using the command line interface (CLI).

Before you begin
Make sure that you have downloaded the CLI client. For more information, see Using the Cloudera Data Engineering command line interface.

About this task

Procedure

1. Run the cde job create command as follows:

```
cde job create --name <JOB_NAME> --type airflow --dag-file <DAG_FILE> --
mount-1-resource <YOUR_DAG_RESOURCE> [--CONFIG-JSON <JSON_STRING> | --
CONFIG-JSON-FILE <FILE_PATH>] [OTHER AIRFLOW FLAGS...]
```

<DAG_file>

is a reference to a file within a Cloudera Data Engineering resource

> **Note:** The --config-json and --config-json-file options allow you to pass configuration parameters to Airflow's trigger DAG run, using them as the conf parameter. These options affect only manually triggered runs and not the scheduled Airflow runs.
>
> If you create a Cloudera Data Engineering Airflow job with a --config-json or --config-json-file option, then a manually triggered Cloudera Data Engineering Airflow job run (through UI, CLI, or REST API) passes these parameters to Airflow. These parameters can be overridden during manual triggers (through CLI or REST API). For more information, see Running an Airflow job using the CLI.

To see the full command syntax and supported options, run cde job create   --help.

> **Note:** Airflow DAGs manage their own schedules and so their schedules cannot be set through the CLI.

You can also hardcode default parameters in the DAG files for scheduled DAG runs. For more information, see Params in the Apache Airflow documentation.

2. Run cde job describe to verify that the job was created:

```
cde job describe --name <JOB_NAME>
```

3. If you want to update the job configuration, use the cde job update command.
For example, to change the number of executors:

```
cde job update --name test_job
```

To see the full command syntax and supported options, run cde job update   --help.

4. To verify the updated configuration, run cde job   describe again:

```
cde job describe --name <JOB_NAME>
```

Listing jobs using the CLI

To view existing applications, run cde job list. To view details for a specific application, run cde job describe --name *<JOB_NAME>*

Submitting a Spark job using the CLI

The following example demonstrates how to submit a JAR or Python file to run on Cloudera Data Engineering Spark in Cloudera Data Engineering using the command line interface (CLI).

About this task

Using the cde spark submit command is a quick and efficient way of testing a spark job, as it spares you the task of creating and uploading resources and job definitions before running the job, and cleaning up after running the job.

This command is recommended only for JAR or Python files that need to be run just once, because the the file is removed from Spark at the end of the run. To manage jobs that need to be run more than once, or that contain schedules, use cde   job run instead of this command.

Procedure

To submit a JAR or Python file to run on Cloudera Data Engineering Spark, use the CLI command:

```
cde spark submit <JAR/Python file> [args...] [Spark flags...] [--job-name <j
ob name>] [--hide-logs]
```

You can use [--job-name <job name>] to specify the same Cloudera Data Engineering job name for consecutive cde spark submit commands. To see the full command syntax and supported options, run cde spark   submit --help.

For example:

To submit a job with a local JAR file:

```
cde spark submit my-spark-app-0.1.0.jar 100 1000 --class com.company.app.spa
rk.Main
```

The CLI displays the job run ID followed by the driver logs, unless you specified the --hide-logs option. The script returns an exit code of 0 for success or 1 for failure.

Running raw Scala code in Cloudera Data Engineering

Cloudera Data Engineering supports running raw Scala code from the command line, without compiling it into a JAR file. You can use the cde spark submit command to run a .scala file. Cloudera Data Engineering recognizes the file as Scala code and runs it using spark-shell in batch mode rather than spark-submit.
Limitations:

- When setting the Log Level from the user interface, the setting is not applied to the raw Scala jobs.
- Do not use package <something> in the raw Scala job file as Raw Scala File is used for Scripting and not for Jar development and packaging.

Note: Cloudera Data Engineering does not currently support interactive sessions. The Scala code runs in batch mode spark-shell.

Run cde spark submit as follows to run a Scala file:

```
cde spark submit filename.scala --jar <JAR_DEPENDENCY_1> --j
ar <JAR_DEPENDENCY_2> ...
```

Submitting an Airflow job using the CLI

The following example demonstrates how to submit a DAG file to immediately run on Cloudera Data Engineering Airflow in Cloudera Data Engineering using the command line interface (CLI).

About this task

Using the cde airflow submit command is a quick and efficient way of testing an Airflow job, as it spares you the task of creating and uploading resources and job definitions before running the job, and cleaning up after running the job.

This command is recommended only for Airflow DAGs that need to be run just once, because the DAG is removed from Airflow at the end of the run. To manage Airflow DAGs that need to be run more than once, or that contain schedules, use cde   job run instead of this command.

Procedure

To submit a DAG file to run on Cloudera Data Engineering Airflow, use the CLI command:

```
cde airflow submit <DAG python file> [--config-json <json-string> | --config
-json-file <file_path>]* [--job-name <job name>]
```

To see the full command syntax and supported options, run cde airflow   submit --help.

For example:

To submit a job with a local DAG file:

```
cde airflow submit my-dag.py
```

When the job has been submitted the CLI displays the job run ID, waits for the job to terminate, and returns an exit code of 0 for success or 1 for failure.

Running a Spark job using the CLI

The following example demonstrates how to run a Cloudera Data Engineering Spark job using the command line interface (CLI).

Before you begin

Make sure that the Spark job has been created and all necessary resources have been created and uploaded.

> **Note:** Custom Docker container images is a *Technical Preview* feature. Contact your Cloudera account representative to enable access to this feature.

About this task

Using the cde job run requires more preparation on the target environment compared to the cde spark submit command. Whereas cde spark submit is a quick and efficient way of testing a Spark job during development, cde job run is suited for production environments where a job is to be run multiple times, therefore removing resources and job definitions after every job run is neither necessary, nor viable.

Procedure

To run a Spark job, run the following command:

```
cde job run --name <job name> [Spark flags...] [--wait] [--variable name=val
ue...]
```

- With [Spark flags...] you can override the corresponding job values. Spark flags that can be repeated replace the original list, except for --conf which only adds or replaces values for the given keys.
- With [--variable] flags you can replace strings in job values. Currently the supported fields are:

  - Spark application name
  - Spark arguments
  - Spark configurations

  For a variable flag name=value any substring {{{name}}} in the value of the supported field gets replaced with value.
- A custom runtime Docker image can be specified for the job using the --runtime-image-resource-name flag, which has to refer to the name of a custom image resource that has already been created.
- GPU Acceleration (Technical Preview): Using [--enable-gpu-acceleration] you can accelerate your Spark jobs using GPUs. You can use [--executor-node-selector   "nvidia.com/gpu=A100"] and [--executor-node-tolerat

ion "nvidia.com/gpu=true"] options to configure selectors and tolerations if you want to run the job on specific GPU nodes. When this job is run, this particular job will request GPU resources.

> ⚠️ **Warning:** You must ensure this virtual cluster has been configured with GPU resource quota. Otherwise, the jobs will be in the Pending state as no GPU resource can be allocated to the pod.

```
cde job run --name example-pi \
--enable-gpu-acceleration \
--executor-node-selector "nvidia.com/gpu=A100" \
--executor-node-toleration "nvidia.com/gpu=true"
```

By default the command returns the job run ID as soon as the job has been submitted.

Optionally, you can use the --wait switch to wait until the job run ends and returns a non-zero exit code if the job run was not successful.

Running an Airflow job using the CLI

The following example demonstrates how to run a Cloudera Data Engineering Airflow job using the command line interface (CLI).

Before you begin

Make sure that the job has been created and all necessary resources have been created and uploaded.

> 📝 **Note:** Custom Docker container images is a *Technical Preview* feature. Contact your Cloudera account representative to enable access to this feature.

About this task

Using the cde job run requires more preparation on the target environment compared to the cde airflow submit command. Whereas cde airflow submit is a quick and efficient way of testing an Airflow job during development, cde job run is suited for production environments where a job is to be run multiple times, therefore removing resources and job definitions after every job run is neither necessary, nor viable.

Procedure

To run an Airflow job, run the following command:

```
cde job run --name <job name> [--config-json <JSON_string> | --config-json-f
ile <file_path>]* [--wait]
```

Airflow configs provided at job run time will override the corresponding job configs.

> 📝 **Note:** The --config-json and --config-json-file options allow you to pass configuration parameters to Airflow's trigger DAG run, using them as the conf parameter. These options affect only manually triggered runs and not the scheduled Airflow runs.
>
> If you create a Cloudera Data Engineering Airflow job with a --config-json or --config-json-file option, then a manually triggered Cloudera Data Engineering Airflow job run (through UI, CLI, or REST API) passes these parameters to Airflow. These parameters can be overridden during manual triggers (through CLI or REST API). For more information, see Running an Airflow job using the CLI.

By default the command returns the job run ID as soon as the job has been submitted.

Optionally, you can use the --wait switch to wait until the job run ends and returns a non-zero exit code if the job run was not successful.

Scheduling Spark jobs

Spark jobs can optionally be scheduled so that they are automatically run on an interval. Cloudera Data Engineering uses the Apache Airflow scheduler to create the schedule instances.

About this task

Note:

Airflow DAGs manage their own schedules, therefore Airflow job schedules cannot be set in this way, other than by using the operational commands pause, unpause, clear, mark-success.

Before you begin

Make sure that the Spark job has been created and all necessary resources have been created and uploaded.

Note: Custom Docker container images is a *Technical Preview* feature. Contact your Cloudera account representative to enable access to this feature.

Procedure

1. Define a running interval for your Spark job:

   The schedule interval is defined by a cron expression. Intervals can be regular, such as daily at 3 a.m., or irregular, such as hourly but only between 2 a.m. and 6 a.m. and only on weekdays. You can provide the cron expression directly or you can generate it using flags.

   > **Note:** Scheduled job runs start at the end of the first full schedule interval after the start date, at the end of the scheduled period. For example, if you schedule a job with a daily interval with a start_date of 14:00, the first scheduled run is triggered at the end of the next day, after 23:59:59. However if the start_date is set to 00:00, it is triggered at the end of the same day, after 23:59:59.

   Available schedule interval flags are:

   --cron-expression

   > A cron expression that is provided directly to the scheduler. For example, 0 */1 * * *

   --every-minutes

   > Running frequency in minutes. Valid values are 0-59. Only a single value is allowed.

   --every-hours

   > Running frequency in hours. Valid values are 0-23. Only a single value is allowed.

   --every-days

   > Running frequency in days. Valid values are 1-31. Only a single value is allowed.

   --every-months

   > Running frequency in months. Valid values are 1-12. Only a single value is allowed.

   --for-minutes-of-hour

   > The minutes of the hour to run on. Valid values are 0-59. Single value, range (e.g.: 1-5), or list (e.g.: 5,10) are allowed.

   --for-hours-of-day

   > The hours of the day to run on. Valid values are 0-23. Single value, range (e.g.: 1-5), or list (e.g.: 5,10) are allowed.

   --for-days-of-month

   > The days of the month to run on. Valid values are 1-31. Single value, range (e.g.: 1-5), or list (e.g.: 5,10) are allowed.

   --for-months-of-year

   > The months of the year to run on. Valid values are 1-12 and JAN-DEC. Single value, range (e.g.: 1-5), or list (e.g.: APR,SEP) are allowed.

   --for-days-of-week

   > The days of the week to run on. Valid values are SUN-SAT and 0-6. Single value, range (e.g.: 1-5), or list (e.g. TUE,THU) are allowed.

   For example, to set the interval as hourly but only between 2 a.m. and 6 a.m. and only on weekdays, use the command:

   ```
   cde job create --name test_job --schedule-enabled=true --every-hours 1 --
   for-minutes-of-hour 0 --for-hours-of-day 2-6 --for-days-of-week MON-FRI --
   schedule-start 2021-03-09T00:00:00Z
   ```

   Or, equivalently, using a single cron expression:

   ```
   cde job create --name test_job --schedule-enabled=true --cron-expression
    '0 2-6/1 * * MON-FRI'  --schedule-start 2021-03-09T00:00:00Z
   ```

2. Define a time range for your Spark job:

The schedule also defines the range of time that instances can be created for. The mandatory --schedule-start flag timestamp tells the scheduler the date and time from which the scheduling begins. The optional --schedule-end flag timestamp tells the scheduler the last date and time at which the schedule is active. If --schedule-end is not specified, the job runs at the scheduled interval until it is stopped manually.

Note: Timestamps must be specified in ISO-8601 UTC format ('yyyy-MM-ddTHH:mm:ssZ'). UTC offsets are not supported.

For example, to create a schedule that runs at midnight for each day of a single week, use the following command:

```
cde job create --name test_job --schedule-enabled=true --every-days 1 --
for-minutes-of-hour 0 --for-hours-of-day 0 --schedule-start 2021-03-09T0
0:00:00Z --schedule-end 2021-03-15T00:00:00Z
```

Enabling, disabling, and pausing scheduled jobs

Using the Cloudera Data Engineering command line interface (CLI), you can enable, disable, or pause scheduled job runs.

Before you begin

Note:

Disabling the schedule removes all record of prior schedule instances.

Note:

Pausing and unpausing the schedule does not remove the record of prior schedule instances.

Procedure

• To enable or disable a job schedule, use the following command:

```
cde job (create | update) --name <job name> --schedule-enabled=(true | f
alse) ...
```

• To pause a job schedule upon schedule creation:

```
cde job (create | update) --name <job name> --schedule-enabled=true --sc
hedule-paused=true ...
```

• To pause an existing job schedule:

```
cde job schedule pause --name <job name>
```

or

```
cde job schedule pause-all
```

• To unpause an existing job schedule:

```
cde job schedule unpause --name <job name>
```

Managing the status of scheduled job instances

Using the Cloudera Data Engineering command line interface (CLI), you can clear the statuses of a range of scheduled instances or mark a scheduled job instance as successful.

Procedure

- To clear the status of a range of scheduled instances, run the following command:

```
cde job schedule clear [--schedule-start <start of clear period>] [--sch
edule-end <end of clear period>]
```

- To mark a single scheduled instance as successful, run the following command:

```
cde job schedule mark-success --execution-date <execution date of schedu
led instance>
```

where <execution date of scheduled instance> is the timestamp that the instance was scheduled for, not when it actually ran.

Managing Sessions in Cloudera Data Engineering using the CLI

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.

Creating a Session using the CDE CLI [Technical Preview]

The cde session create command allows you to create a new Session.

Before you begin

Ensure that you are using a version of Cloudera Data Engineering 1.19 or higher for your Virtual Cluster and ensure that you have downloaded the CLI client. For more information, see Using the Cloudera Data Engineering command line interface.

Procedure

Run the following command in the CDE CLI:

```
cde session create --name <session-name> --type <pyspark/spark-scala>
```

Interacting with a Session using the CDE CLI

Once your Session has been created, you can interact with it using the cde sessions interact command.

About this task

Below is an example that demonstrates how to interact with a PySpark or Scala Session in Cloudera Data Engineering using the CLI.

Before you begin

Ensure that you are using a version of Cloudera Data Engineering 1.19 or higher for your Virtual Cluster and ensure that you have set up the CDE CLI client and have created a Cloudera Data Engineering Session using the CLI or UI. For more information, see Using the Cloudera Data Engineering command line interface.

Procedure

Run the following command in the CDE CLI:

```
cde session interact --name <session-name>
```

Sessions example for the CDE CLI

In this example, a Session is created using the Cloudera Data Engineering CLI with resources specified during creation. In this example, python environment, files, Git repository, and workload credentials resources are used.

Note: Access files from the /app/mount after they are mounted. Access Secrets from /etc/dex/secrets/<secret -name>.

```
> cde session create --name resources --type pyspark --python-env-resource-n
ame example-virtual-env --runtime-image-resource-name docker-image --mount-1
-resource octocat --mount-2-resource example-files --mount-3-resource exampl
e-data --workload-credential workload-cred --workload-credential workload-cr
ed-2
{
  "name": "resources",
  "type": "pyspark",
  "creator": "csso_surya.balakrishnan",
  "created": "2023-10-06T03:13:03Z",
  "mounts": [
    {
      "dirPrefix": "/",
      "resourceName": "octocat"
    },
    {
      "dirPrefix": "/",
      "resourceName": "example-files"
    },
    {
      "dirPrefix": "/",
      "resourceName": "example-data"
    }
  ],
  "lastStateUpdated": "2023-10-06T03:13:03Z",
  "state": "starting",
  "interactiveSpark": {
    "id": 1,
    "driverCores": 1,
    "executorCores": 1,
    "driverMemory": "1g",
    "executorMemory": "1g",
    "numExecutors": 1,
    "pythonEnvResourceName": "example-virtual-env"
  },
  "workloadCredentials": [
    "workload-cred",
    "workload-cred-2"
  ],
  "runtimeImageResourceName": "docker-image"
}

> ./cde session interact --name resources
Starting REPL...
Waiting for the session to go into an available state...
Connected to Cloudera Data Engineering...
Press Ctrl+D (i.e. EOF) to exit
```
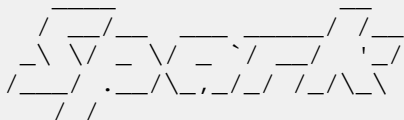
```
Welcome to
      ____              __
     / __/__  ___ _____/ /__
    _\ \/ _ \/ _ `/ __/  '_/
   /___/ .__/\_,_/_/ /_/\_\
      /_/
Type in expressions to have them evaluated.

>>> os.listdir("/app/mount")
['.git', 'README', 'access-logs-ETL-iceberg.py', 'access-logs-ETL.py', 'a
ccess-logs.txt', 'cdeoperator.py', 'pyspark-batch-job.py', 'pyspark_wordcoun
t.py', 'spark-load-data.py', 'word_count_templates.txt', 'wordcount_input_1.
txt', 'dex-spark-driver-template-txckdpxp.yaml', 'dex-spark-executor-templat
e-txckdpxp.yaml']

>>> sec_path = "/etc/dex/secrets/workload-cred/key1"
>>> with open(sec_path) as f:
...     for line in f:
...         print(line)
value1
>>> sec_path = "/etc/dex/secrets/workload-cred-2/key2"
>>> with open(sec_path) as f:
...     for line in f:
...         print(line)
value2
>>> import pandas

>>> dates = pandas.date_range("20130101", periods=6)

>>> dates
DatetimeIndex(['2013-01-01', '2013-01-02', '2013-01-03', '2013-01-04',
               '2013-01-05', '2013-01-06'],
              dtype='datetime64[ns]', freq='D')
```

Sessions command descriptions

The Cloudera Data Engineering command reference is shown below.

Command	Definition
cde session create	Creates a Cloudera Data Engineering session. Sessions are identified by a user-specified name. Sessions have a type that defines the engine that the Session will run on. The 'spark-scala'[Scala REPL] and 'pyspark'[Python REPL] types are currently supported.
cde session interact	Connects to a running session in a Spark shell similar to the interface and submit statements.
cde session kill	Ends a session. The Spark driver and executor processes are stopped.
cde session delete	Deletes a session and removes all references to the session. Logs will no longer be accessible.
cde session list	Lists all sessions. The --output flag can be used to control the output format. You can use the `--output string output format ("table" or "json") (default "table")` flag to specify whether the Session's output must be in a table or JSON format
cde session statements	Lists session statements. The --output flag can be used to control the output format.
cde session describe –name <session_name>	Describes the session. The command is used as an input to name the session.

Suspending and resuming Cloudera Data Engineering Services using the CLI - Technical preview

This document provides a step-by-step guide on how to suspend and resume Cloudera Data Engineering Services using the CLI, enabling efficient management of resources and cost optimization.

About this task

In AWS

Suspending the Service scales down all compute nodes used by the Cloudera Data Engineering Service and its Virtual Clusters (VCs) to zero.

In Azure

The feature stops the AKS cluster during suspension and restarts it when the Service is resumed.

Before you begin

To suspend and resume your Cloudera Data Engineering Services, you need to:

- Request the DE_ENABLE_SUSPEND_RESUME entitlement
- Have DE Admin user privileges at the environment level
- Make sure that you are using the CDP CLI Beta version 0.9.141 because the suspend-resume-service CLI option is only available in it. To install CDP CLI Beta version, see Installing Beta CDP CLI.

Procedure

1. Suspend your Cloudera Data Engineering Service.

a) Enter maintenance mode.

The maintenance mode stops all running jobs and pauses all scheduled jobs.

```
cdp de suspend-resume-service --cluster-id <value> --step "PREPARE"
```

Once the maintenance step completes, you can suspend the Service.

b) Suspend the Cloudera Data Engineering Service.

Depending on the cloud provider, this operation scales down the cluster to zero (AWS) or shuts down the cluster (Azure).

```
cdp de suspend-resume-service --cluster-id <value> --step "SUSPEND"
```

c) Check the suspend-resume status of the Virtual Cluster.

```
cdp de get-suspend-resume-status --cluster-id <value>
```

2. Resume the Cloudera Data Engineering Service.

   a) Resume your suspended Cloudera Data Engineering Service.

   This action scales up the compute nodes.

   ```
   cdp de suspend-resume-service --cluster-id <value> --step "RESUME"
   ```

   Once the Service and Virtual Clusters are all running, you can exit maintenance mode.

   b) Check the suspend-resume status of the Virtual Cluster.

   ```
   cdp de get-suspend-resume-status --cluster-id <value>
   ```

   c) Exit maintenance mode.

   This action restores the Cloudera Data Engineering cluster to an operational state, and resumes all previously paused scheduled jobs.

   ```
   cdp de suspend-resume-service --cluster-id <value> --step "QUIT"
   ```

3. If the suspend or resume operation fails, try once again before raising a support ticket.

   For more information, see Troubleshooting failures in suspend and resume.

Related Information

Suspending and resuming a Cloudera Data Engineering Service using a CLI script

Suspending and resuming Cloudera Data Engineering Services

Suspending and resuming Cloudera Data Engineering Services using the API

Troubleshooting failures in suspend and resume

Suspending and resuming a Cloudera Data Engineering Service using a CLI script - Technical preview

You have the option to leverage the suspend and resume script to suspend and resume your Cloudera Data Engineering Services.

About this task

After specifying the cluster_id, the operation type, and the profile, the suspend-resume.py script suspends or resumes the entire cluster that is specified by the cluster_id.

Before you begin

To suspend and resume your Cloudera Data Engineering Services, you need to:

- Request the DE_ENABLE_SUSPEND_RESUME entitlement
- Have DE Admin user privileges at the environment level

Procedure

1. Download the suspend-resume.py script from GitHub.

2. Using the CDP CLI, run the suspend and resume script:

```
suspend-resume.py [-h] [--cluster_id [***CLUSTER-ID***]] [--operation [*
**OPERATION***]] [--profile [***PROFILE***]]
```

Specify the following parameters:

- The cluster ID
- The operation type (suspend or resume)
- The profile name

3. If the suspend or resume operation fails, try once again before raising a support ticket.

For more information, see Troubleshooting failures in suspend and resume.

Related Information

Suspending and resuming Cloudera Data Engineering Services

Suspending and resuming Cloudera Data Engineering Services using the CLI

Suspending and resuming Cloudera Data Engineering Services using the API

Troubleshooting failures in suspend and resume

Suspending and resuming Cloudera Data Engineering virtual clusters using the CDP CLI [Technical Preview]

Learn about how to suspend and resume Cloudera Data Engineering Virtual Clusters (VCs) using the CDP CLI.
The Virtual Cluster-level suspend and resume feature enables you to temporarily pause one or more Cloudera
Data Engineering VCs during idle periods. Cloudera has observed that suspending a larger number of VCs offers
significant cost savings, primarily on the compute resources.

About this task

Note: This feature is in Technical Preview and is not ready for production deployments. Cloudera
recommends trying this feature in test or development environments and encourages you to provide feedback
on your experiences.

The virtual cluster-level suspend and resume feature is supported on both AWS and Azure with no operational
differences on these two cloud providers.

The virtual cluster-level suspend and resume feature is implemented as a convenient, user-friendly, and one-step
procedure.

When you suspend a VC

1. The VC is transitioned to maintenance mode, where:

   - The VC is locked
   - Running jobs and sessions are killed
   - Job schedules are paused

2. The suspend operation is started by scaling down all the infrastructure deployments of the given
   VC.

   Depending on the distribution of the pods, the associated node can be released.

When you resume a VC

1. First the VC is resumed.

    2. The VC exits maintenance mode, which involves:

- Unpausing all the job schedules
- Unlocking the VC

Before you begin

To suspend and resume your Cloudera Data Engineering virtual cluster, you need to:

- Request the DE_ENABLE_VC_SUSPEND_RESUME entitlement
- Have DEAdmin user privileges at the environment level
- Set up ~/.cdp/config
- Set up ~/.cdp/credentials
- Make sure that you are using the CDP CLI Beta version because the suspend-vc and the resume-vc CLI options are only available in it. To install the CDP CLI Beta version, see Installing Beta CDP CLI.

Procedure

1. Install and run the BETA CDP CLI.

2. Suspend your Cloudera Data Engineering virtual cluster.

```
cdp de suspend-vc --cluster-id <cluster-xxxxxx> --vc-id <dex-app-xxxxxx>
```

cluster-id

> To obtain the cluster ID, log in to the Cloudera Data Engineering UI and navigate to Administration in the left-navigation menu. Select the respective service, then click on the Service Details icon, and copy the CLUSTER ID.

vc-id

> To obtain the VC ID, log in to the Cloudera Data Engineering UI and navigate to Administration in the left-navigation menu. Select the respective service and then the respective VC. Click on the Virtual Cluster Details icon, and copy the VC-ID.

3. Resume the previously suspended Cloudera Data Engineering virtual cluster.

```
cdp de resume-vc --cluster-id <cluster-xxxxxx> --vc-id <dex-app-xxxxxx>
```

4. Check the suspend-resume status of the virtual cluster.

```
cdp de get-suspend-resume-status --cluster-id <cluster-xxxxxx> --vc-id <
dex-app-xxxxxx>
```

> **Note:** If you do not pass the --vc-id flag in the command, the CDP CLI call returns the suspend-resume status for the Cloudera Data Engineering service.

Related Information

Overview of suspending and resuming Cloudera Data Engineering virtual clusters [Technical Preview]

Suspending and resuming Cloudera Data Engineering virtual clusters using the API [Technical Preview]

Creating a Git repository in Cloudera Data Engineering using the CLI

Learn how to create a Git repository in Cloudera Data Engineering using the CLI with the supported Git providers, including GitHub, GitLab, and Bitbucket.

Before you begin

For information on limitations and how to create credentials, see Creating a Git repository in Cloudera Data Engineering.

Procedure

- Using the CDE CLI, run the following command:

```
cde repository create --name [***REPOSITORY-NAME***] --url [***REPOSITORY-
URL***] --branch [***BRANCH-NAME***] --insecure-skip-tls --credential [***
USER-CREDENTIAL***]
```

Provide the values for the following parameters:

--name

The repository name.

--url

The repository URL.

--branch

The branch name.

--insecure-skip-tls

Optional parameter. Use this parameter if the server uses a self-signed CA certificate that Cloudera Data Engineering does not trust. This allows Cloudera Data Engineering to skip the security check and clone the repository.

--credential

If the repository does not require credentials, then this parameter is optional. For information on how to create credentials, see Creating a Git repository in Cloudera Data Engineering.

Cloudera Data Engineering Spark job example

In this example there is a local Spark jar my-app-0.1.0.jar, and a local reference file my-ref.conf that the Spark job opens locally as part of its execution. The Spark job reads data from the location in the first argument and writes data to the location in the second argument. There is also a custom Spark configuration for tuning performance.

1. Make your job available for running in one of the following ways:

You can submit the job directly to Cloudera Data Engineering and have it run the job once, using the spark submit command. In this case no permanent resources are created on Cloudera Data Engineering subsequently no cleanup is necessary after the job run. This is ideal when testing a job.

```
cde spark submit my-app-0.1.0.jar \
  --file my-ref.conf \
  --conf spark.sql.shuffle.partitions=1000
```

If you plan to run the same job several times it is is a good idea to create and upload the resource and job and then run it on Cloudera Data Engineering using the job   run command. This is the preferable method in production environments.

```
> cde resource create --name my-resource
> cde resource upload --name my-resource --local-path my-app-0.1.0.jar
 109.7MB/109.7MB 100% [=========================================] my-
app-0.1.0.jar
> cde resource upload --name my-resource --local-path my-ref.conf
```

```
    135.0b/135.0b 100% [=============================================] my-
ref.conf
> cde job create \
   --name my-job \
   --type spark \
   --mount-1-resource my-resource \
   --application-file my-app-0.1.0.jar \
   --conf spark.sql.shuffle.partitions=1000 \
> cde job run --name my-job
{
   "id": 1
}
> cde run describe --id 1 | jq -r '.status'
starting
...
> cde run describe --id 1 | jq -r '.status'
finished
```

2. Schedule your job:

As the above created job stays in Cloudera Data Engineering permanently until you delete it, you can schedule it to run regularly at a predefined time. This example schedules your job to run daily at midnight, starting from January 1, 2021:

```
> cde job update \
   --name my-job \
   --schedule-enabled=true \
   --schedule-start 2021-01-01T00:00:00Z \
   --every-days 1 \
   --for-minutes-of-hour 0 \
   --for-hours-of-day 0
```

CDE CLI command reference

The Cloudera Data Engineering command line syntax is shown below. You can view additional syntax help by adding --help after any command.

cde command

```
Usage:
  cde [command]

Available Commands:
  airflow     Airflow commands
  backup      Create and Restore CDE backups
  credential  Manage CDE credentials
  help        Help about any command
  job         Manage CDE jobs
  profile     Manage CDE configuration profiles
  repository  Manage CDE repository
  resource    Manage CDE resources
  run         Manage CDE runs
  session     Manage CDE sessions
  spark       Spark commands
Flags:
      --access-key-id string        access key identifier
      --access-key-secret string    access key secret
      --auth-cache-file string      token file cache location (default "$
USERCACHE/token-cache")
      --auth-no-cache               do not cache authentication tokens
```

```
      --auth-pass-file string      authentication password file location
      --cdp-endpoint string        CDP API endpoint (default depends on
 CDP Control Plane region)
      --config-profile string      CDE configuration profile name
      --credentials-file string    CDP credentials file location
      --credentials-profile string  CDP credentials profile name (default
 "default")
  -h, --help                       help for cde
      --hide-progress-bars         hide progress bars for file uploads
      --region string              CDP Control Plane region ("us-west-1",
 "eu-1" or "ap-1") (default "us-west-1")
      --skip-credentials-file      skip CDP credentials file discovery
      --tls-ca-certs string        additional PEM-encoded CA certificates
      --tls-insecure               skip verification of API server TLS cer
tificate
      --user string                CDP user to authenticate as
      --vcluster-endpoint string   CDE virtual cluster endpoint
  -v, --verbose                    verbose logging
      --version                    version for cde

Use "cde [command] --help" for more information about a command.
```

cde airflow command

```
Usage:
  cde airflow [command]

Available Commands:
  activate-pyenv   Deprecated, please use /admin/airflow/env APIs or UI inst
ead. Can be used ONLY for the private cloud setup. Activate an existing envi
ronment for a given VC
  delete-dag       Delete an unassociated DAG from Airflow
  get-active-pyenv Deprecated, please use /admin/airflow/env APIs or UI in
stead. Can be used ONLY for the private cloud setup. Get currently active ai
rflow python env resource name
  reset-pyenv      Deprecated, please use /admin/airflow/env APIs or UI i
nstead. Can be used ONLY for the private cloud setup. Reset Airflow env to f
actory defaults
  submit           Run a DAG on CDE Airflow
```

cde backup command

```
Usage:
  cde backup [command]

Available Commands:
  create       Create a backup of jobs and resources
  list-archive List jobs and resources in an archive
  restore      Restore a backup of jobs and resources
```

cde credential command

```
Usage:
  cde credential [command]

Available Commands:
  create     Create a credential (experimental)
  delete     Delete a credential (experimental)
  describe   Describe a credential (experimental)
  list       List credentials (experimental)
```

```
  update      Update a credential (experimental)
```

cde job command

```
Usage:
  cde job [command]

Available Commands:
  clone       Clone an existing job or job run as a new job
  create      Create a job
  delete      Delete a job
  describe    Describe a job
  import      Import a job
  list        List jobs
  run         Run a job
  schedule    Operate CDE job schedules
  update      Update a job
```

cde profile command

```
Usage:
  cde profile [command]

Available Commands:
  show-active Shows the active profile of config file
```

cde repository command

```
Usage:
  cde repository [command]

Available Commands:
  create      Create a repository
  delete      Delete a repository
  describe    Describe repository
  download    Download a file from a repository
  list        List repositories
  sync        Syncs a repository with its remote
```

cde resource command

```
Usage:
  cde resource [command]
Available Commands:
  create           Create a resource
  delete           Delete a resource
  delete-file      Delete a file from a resource
  describe         Describe resource
  download         Download a file from a resource
  download-archive Download resource as archive
  list             List resources
  list-events      List resource events
  upload           Upload one or more files to a resource
  upload-archive   Upload and extract archive to resource
```

cde run command

```
Usage:
  cde run [command]

Available Commands:
  debug        Debug a run
  describe     Describe a run
  fg-status    Get the fine-grained statuses of a run
  kill         Kill a run
  list         List runs
  logs         Retrieve logs for a run
  ui           Open a run in the default browser
  wait         Wait for a run to end
```

cde session command

```
Usage:
  cde session [command]

Available Commands:
  create       Create a session
  delete       Delete a session
  describe     Describe a sessions
  interact     interact with a running session
  kill         Kill a session
  list         List sessions
  logs         Retrieve logs for a session
  statements   List session statements
```

cde spark command

```
Usage:
  cde spark [command]

Available Commands:
  submit        Run a jar/py file on CDE Spark
```

CDE CLI Spark flag reference

The Cloudera Data Engineering command Spark flag reference is shown below.

```
--application-file: application main file
--class: application main class
--arg: Spark argument
--conf: Spark configuration (format key=value) (can be repeated)
--min-executors: minimum number of executors
--max-executors: maximum number of executors
--initial-executors: initial number of executors
--executor-cores: number of cores per executor
--executor-memory: memory per executor
--driver-memory: memory for driver
--driver-cores: number of driver cores
--spark-name: Spark application name
--file: additional file additional file (can be repeated) (will be merged w
ith --files, if provided)
--files: additional files (comma-separated list) (will be merged with all --
file)
```

```
--jar: additional jar (can be repeated) (will be merged with --jars, if prov
ided)
--jars: additional jars (comma-separated list) (will be merged with all --
jar)
--py-file: additional Python file (can be repeated) (will be merged with --
py-file, if provided)
--py-files: additional Python files (comma-separated list) (will be merged
 with all --py-file)
--packages: additional dependencies as comma-separated list of Maven coordi
nates
--repositories: additional repositories/resolvers for retrieving the --pac
kages dependencies
--python-env-resource-name: Python environment resource name
--python-version: Python version ("python3" or "python2")
--log-level: log level for Spark containers (TRACE, DEBUG, INFO, WARN, ERR
OR, FATAL, OFF)
--enable-analysis: enables Spark analysis (see 'Analysis' UI tab for a job r
un)
```

CDE CLI Airflow flag reference

The Cloudera Data Engineering command Airflow flag reference is shown below.

```
cde airflow submit --help
Usage:
  cde airflow submit [flags]

Examples:
For a local DAG file 'my-airflow-job.py':
> cde airflow submit my-airflow-job.py


Flags:
      --airflow-file-mount-N-prefix string     mount directory prefix for
 airflow file mount N (defaults to "/airflow-file-mount-N-resource-name")
      --airflow-file-mount-N-resource string   resource name for airflow f
ile mount N
      --config stringArray                     DEPRECATED - DAG configuratio
n for manual triggers (format key=value) (can be repeated). Use --config-jso
n option instead.
      --config-json string                     DAG configuration for manual
 triggers in JSON string format
      --config-json-file string                DAG configuration file locat
ion for manual triggers in JSON format
      --dag-file string                        DAG filename, path to the
DAG within the resource
  -h, --help                                   help for submit
      --job-name string                        name of the generated job
```

CDE CLI list command syntax reference

You can include flags with the Cloudera Data Engineering command line interface (CLI) list command calls to filter the result set.

cde [credential|job|resource|run|...] list [--filter   [fieldname[operator]argument]] [--filter [fieldname[operator]argume
nt]] ...

A list command call can include multiple filter flags, where all filters must match for the entry to be returned. You
have to enclose filters in quotes.

fieldname

> is selected from the top-level fields of the returned entries. Filtering of fields nested within other
> fields is supported using MySQL 8 JSON path expressions.

operator

> is one of: eq, noteq, lte, lt, gte, gt, in, notin, like, rlike. The in and notin operators work on an
> argument of comma-separated values. The like operator matches using SQL LIKE syntax, e.g. %test
> %. The rlike operator matches using the SQL REGEXP regular expression syntax.

argument

> is the value, list, or expression to match with the operator. If the argument contains commas the
> filter has to be enclosed in a second set of quotes, for example: '"id[in]12,14,16"'.

> **Note:**
>
> Timestamps must be formatted as MySQL date time literals.

For example:

```
cde run list --filter 'spark.spec.file[rlike]jar'
```

Managing workload secrets with Cloudera Data Engineering Spark Jobs using the CLI

Cloudera Data Engineering provides a secure way to create and store workload secrets for Cloudera Data Engineering
Spark Jobs. This is a more secure alternative to storing credentials in plain text embedded in your application or job
configuration.

Creating a workload secret for Cloudera Data Engineering Spark Jobs using CLI

You can create workload secrets using interactive mode or using a JSON file. You can use the `--workload-cred-json-file` and the `--workload-cred-key` flags along with the `--type` flag supporting workload-credential.

Before you begin

Make sure that you have downloaded the CLI client. For more information, see Using the Cloudera Data Engineering
command line interface.

For Interactive mode

- Specify --workload-cred-key when prompted for secret values multiple times. The values which are sensitive are read as a hidden password field interactively from the CLI.

```
./cde credential create --name <workload-credential-name> --type workloa
d-credential --workload-cred-key <workload-credential_key> --workload-cr
ed-key <workload-credential_key>
```

For example:

```
./cde credential create --name workload-cred-1 --type workload-credentia
l --workload-cred-key db-pass --workload-cred-key aws-secret

Enter Secret value for Workload Cred key "db-pass" :
Re-enter Secret value for Workload Cred key "db-pass" :
Enter Secret value for Workload Cred key "aws-secret" :
Re-enter Secret value for Workload Cred key "aws-secret" :
```

For JSON file

1. Create a JSON file with workload secret keys.

```
sample.json file

{
    "<workload-credential-key>": "<secret_value_of_key>",
    "<workload-credential-key>": "<secret_value_of_key>"
  }
```

For example:

```
sample.json file

{
    "aws-secret": "secret123",
    "db-pass": "dbpass123"
  }
```

2. Run the following command to create the workload secret:

```
./cde credential create  --name <workload-credential-name> --type worklo
ad-credential --workload-cred-json-file <workload-credential-json-file-n
ame>
```

For example:

```
./cde credential create  --name workload-cred-1 --type workload-credenti
al --workload-cred-json-file sample.json
```

Updating a workload secret for Cloudera Data Engineering Spark Jobs using CLI

You can update an existing secret to use it with the Cloudera Data Engineering Spark Jobs.

For Interactive mode

- Specify --workload-cred-key when prompted for secret values multiple times. The values which are sensitive
  are read as a hidden password field interactively from the CLI.

```
./cde credential update --name <workload-credential-name> --workload-cre
d-key <workload-credential_key> --workload-cred-key <workload-credential
_key>
```

For example:

```
./cde credential update --name workload-cred-1 --workload-cred-key db-pa
ss --workload-cred-key aws-secret --workload-cred-key api-token
Enter Secret value for Workload Cred key "dbPassword" :
Re-enter Secret value for Workload Cred key "dbPassword" :
Enter Secret value for Workload Cred key "aws-secret" :
Re-enter Secret value for Workload Cred key "aws-secret" :
```

For JSON file

1. Update the JSON file with workload secret keys.

   For example:

```
sample.json file – file name

{
    "aws-secret": "secret123",
    "db-pass": "dbpass123"
  }
```

2. Run the following command to create the workload secret with updated parameters:

```
./cde credential update --name <workload-credential-name> --type workloa
d-credential --workload-cred-json-file <workload-credential-json-file-na
me>
```

For example:

```
./cde credential update --name workload-cred-2 --workload-cred-json-file
 sample.json
```

Linking a workload secret to the Cloudera Data Engineering Spark Job definitions using CLI

After you create a workload secret, you can link it to the Cloudera Data Engineering Spark Job definitions that you
created using CLI.

```
./cde job create --name <workload-credential-name> --type <workload-credenti
al_type> --application-file <application-file-name> --workload-credential <w
orkload-credential-1> --workload-credential <workload-credential-2>
```

For example:

```
./cde job create --name test-workload-job --type spark --mount-1-resource te
st-workload --application-file test-workload-cred.py --workload-credential w
orkload-cred-1 --workload-credential workload-cred-2
```

Using a workload secret in Spark application code

To use the workload secret credentials, you can read the credentials that are mounted into the Spark drivers and executors as read-only files.

The workload secrets are mounted into the Spark drivers and executors in this path:

```
/etc/dex/secrets/<workload-credential-name>/<credential-key-1> /etc/dex/secr
ets/<workload-credential-name>/<credential-key-2>
```

Example workload credentials to use in the application code:

The workload credential is created with the command below.

```
./cde credential create --name workload-cred-1 --type workload-credential --
workload-cred-key db-pass --workload-cred-key aws-secret
```

The secrets can be read as local files from the paths below within the Spark drivers and executors:

```
/etc/dex/secrets/workload-cred-1/aws-secret
/etc/dex/secrets/workload-cred-1/db-pass
```

Example of a PySpark application code to read a secret:

```
from pyspark.sql import SparkSession

spark = SparkSession \
    .builder \
    .appName("Sample DB Connection") \
    .getOrCreate()

# read the password from the local file
dbPass=open("/etc/dex/secrets/workload-cred-1/db-pass").read()

# use the password in a jdbc connection
jdbcDF= spark.read \
        .jdbc("jdbc:postgresql:dbserver", "schema.tablename",
        properties={"user": "username", "password": dbPass})
```

Listing an existing workload secret to the Cloudera Data Engineering Spark Job

You can list an existing secret for Cloudera Data Engineering Spark Jobs using CLI.

```
./cde credential list  --filter 'type[eq]workload-credential'
```

Example output:

```
[
    {
      "name": "workload-cred-1",
      "type": "workload-credential",
      "description": "workload credential description",
      "created": "2022-12-18T07:26:41Z",
      "modified": "2022-12-18T07:26:41Z"
    }
  ]
```

Deleting a workload secret for Cloudera Data Engineering Spark Jobs using CLI

You can delete an existing secret for Cloudera Data Engineering Spark Jobs using CLI.

About this task

```
./cde credential delete --name <workload-credential-name>
```

For example:

```
./cde credential delete --name workload-cred-1
```