

Backing up and restoring Cloudera Data Engineering jobs

Date published: 2020-07-30

Date modified: 2022-11-18

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 a horizontal bar extending to the right.

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

Backing up Cloudera Data Engineering jobs.....	4
Restoring Cloudera Data Engineering jobs from backup.....	6

Backing up Cloudera Data Engineering jobs

You can back up Cloudera Data Engineering (CDE) jobs and associated resources. Backups are saved as zip files that can be used to restore jobs and their associated resources.



Note: Currently, the user who runs the `cde backup restore` command has permissions, by default, to run the Jobs. This may cause CDE jobs to fail if the workload user differs from the user who runs the Jobs on Source CDE Service where the backup was performed. This failure is due to the Workload User having different privileges as the user who is expected to run the job.

For CDE CLI

Before you begin

- Download and configure the CDE CLI.

Steps

1. Run the `cde backup create` command to create a backup of the jobs in the virtual cluster your CDE CLI is configured to interact with. By default, all job configurations in the cluster are backed up, but the resources are not. You can use command flags to change this behavior as follows:

--include-job-resources

Backs up all resources associated with the selected jobs. These resources cannot be filtered out by the `--resource-filter` parameter.

--include-jobs

Backs up all jobs in the virtual cluster. This is the default behavior.

--include-resources

Backs up all resources in the virtual cluster, including those not associated with jobs.

--job-filter *<filter_string>*

Selects jobs to back up using the `fieldname[operator]argument` syntax. This command flag can be repeated.

--resource-filter *<filter_string>*

Selects resources to back up using the `fieldname[operator]argument` syntax. This command flag can be repeated.

--local-path *<filename>*

Specifies the local file path and name to store the backup. By default, the backup filename is `archive-<timestamp>.zip`.

For example, to backup all jobs containing the string `etl` and include all resources associated with those jobs:

```
cde backup create --job-filter "name[like]%etl%" --include-job-resources
```

2. Validate the backup using the `cde backup list-archive` command. For example:

```
cde backup list-archive --local-path archive-2021-11-10T01:24:06.zip
```

Confirm that all jobs and resources that you expected to be backed up are included.

Result

The output of `cde backup list-archive` is similar to the following:

```
[
  {
    "entityType": "Resource",
```

```

    "cdeVersion": "1.13.0-b199",
    "clusterID": "cluster-n5kmgmwr",
    "appID": "dex-app-47m2j5qz",
    "name": "etl-jars",
    "adjustedName": "etl-jars",
    "archiveDirectoryPath": "v1/resources/v1/709c2910-9695-4496-8f18-5d39
a8771c7d/",
    "user": "psherman",
    "backupDate": "2021-11-10T01:24:06.34440405Z"
  },
  {
    "entityType": "Job",
    "cdeVersion": "1.13.0-b199",
    "clusterID": "cluster-n5kmgmwr",
    "appID": "dex-app-47m2j5qz",
    "name": "etl-demo",
    "adjustedName": "etl-demo",
    "archiveDirectoryPath": "v1/jobs/v1/507961fd-f787-4b7c-87d7-7b441bbad
e73/",
    "user": "psherman",
    "backupDate": "2021-11-10T01:24:06.34440405Z"
  }
]

```

For CDE API



Note: Some of the command examples provided utilize [jq](#) to parse the JSON output and make it more readable.

Before you begin

Request an access token and save it as an environment variable to use in API calls. For instructions, see [Getting a Cloudera Data Engineering API access token](#).

Steps

1. Determine the API URL for the virtual cluster containing the job you want to back up:
 - a. Navigate to the Cloudera Data Engineering Overview page.
 - b. In the CDE Services column, select the service containing the virtual cluster with the jobs you want to back up.
 - c. In the Virtual Clusters column on the right, click the Cluster Details icon for the virtual cluster containing the jobs you want to back up.
 - d. Copy the URL under JOBS API URL, and set it as an environment variable. For example:

```
export CDE_JOBS_API="https://pmjkrng5.cde-cz1mkz4y.na-01.xvp2-7p8o.c
loudera.site/dex/api/v1"
```

2. Back up jobs using a URL-encoded filter with the syntax `name[like]<query>`, modeled after the SQL like operator. For example, to back up jobs containing the string `etl`, set `jobfilter` to `name[like]etl%` (URL-encoded as `name%5Blike%5D%25etl%25`):

```
curl -k \
-H "Authorization: Bearer ${CDE_TOKEN}" \
-X GET "${CDE_JOBS_API}/admin/export?exportjobs=true&jobfilter=name%5Bli
ke%5D%25etl%25&exportjobresources=true&exportresources=false" \
-H "accept: application/zip" \
--output cde-etl-jobs-backup.zip
```

To back up all jobs and associated resources, omit the `jobfilter` parameter:

```
curl -k \
```

```
-H "Authorization: Bearer ${CDE_TOKEN}" \
-X GET "${CDE_JOBS_API}/admin/export?exportjobs=true&exportjobresources=
true&exportresources=false" \
-H "accept: application/zip" \
--output cde-all-jobs-backup.zip
```

3. (Optional) You can validate a backup file by uploading it to the /admin/list-archive endpoint. For example, for a backup file named cde-all-jobs-backup.zip:

```
curl -k \
-H "Authorization: Bearer ${CDE_TOKEN}" \
-X POST "${CDE_JOBS_API}/admin/list-archive" \
-H "accept: application/json" \
-H "Content-Type: multipart/form-data" \
-F "file=@/path/to/cde-all-jobs-backup.zip;type=application/zip" \
| jq
```

For Web UI

Before you begin

Steps

1. In the Cloudera Data Platform (CDP) management console, click the Data Engineering tile and click Overview.
2. In the CDE Services column, select the service containing the virtual cluster with the jobs you want to back up.
3. In the Virtual Clusters column on the right, click the View Jobs icon for the virtual cluster with the jobs you want to back up.
4. Click Jobs in the left menu.
5. Click the vertical ellipses menu at the top right, and then click Backup Jobs.

Result

Depending on your browser settings, you are either prompted for a location to save the file, or the file is downloaded to your default download location. The file is a ZIP file named archive-*<timestamp>*.zip.

To restore a backup file, see [Restoring Cloudera Data Engineering jobs from backup](#).

Restoring Cloudera Data Engineering jobs from backup

You can restore Cloudera Data Engineering (CDE) jobs and associated resources from a backup ZIP file.

For CDE CLI

Before you begin

- You must have a valid backup file to restore from. For instructions on backing up CDE jobs, see [Backing up Cloudera Data Engineering jobs](#).
- Download and configure the CDE CLI.

Steps

1. Run the `cde backup restore` command to restore a backup file to the virtual cluster your CDE CLI is configured to interact with. Use the `--duplicate-handling` flag to select the policy for handling duplicate jobs. The possible values are as follows:

error

Return an error if there are duplicate job names, and abort the restore operation. This is the default behavior.

rename

If a job name in the backup conflicts with an existing job, keep the existing job and rename the restored job by appending a numerical identifier to the job name.

keep-original

If a job name in the backup conflicts with an existing job, keep the existing job and do not restore the backed up job.

For example:

```
cde backup restore --local-path archive-2021-11-10T01:24:06.zip --duplicate-handling rename
```

Result

Validate that the jobs and resources were restored by running `cde job list` and `cde resource list`.

For CDE API

Note: Some of the command examples provided utilize `jq` to parse the JSON output and make it more readable.

Before you begin

- You must have a valid backup file to restore from. For instructions on backing up CDE jobs, see [Backing up Cloudera Data Engineering jobs](#).
- Request an access token and save it as an environment variable to use in API calls. For instructions, see [Getting a Cloudera Data Engineering API access token](#).

Steps

1. Determine the API URL for the virtual cluster that you want to restore the jobs and resources to:
 - a. Navigate to the Cloudera Data Engineering Overview page.
 - b. In the CDE Services column, select the service containing the virtual cluster where you want to restore the jobs.
 - c. In the Virtual Clusters column on the right, click the Cluster Details icon for the virtual cluster you want to restore to.
 - d. Copy the URL under JOBS API URL, and set it as an environment variable. For example:

```
export CDE_JOBS_API="https://pmjkrn5.cde-czlmkz4y.na-01.xvp2-7p8o.cloudera.site/dex/api/v1"
```

2. Restore jobs from the backup file by uploading the backup file to the `/admin/import` endpoint. You can choose how to handle duplicate job names using the `duplicatehandling=<policy>` parameter. Options are:

error

Return an error if there are duplicate job names, and abort the restore operation.

rename

If a job name in the backup conflicts with an existing job, keep the existing job and rename the restored job by appending a numerical identifier to the job name.

keep-original

If a job name in the backup conflicts with an existing job, keep the existing job and do not restore the backed up job.

For example, to restore a backup named `cde-etl-jobs-backup.zip` using the `rename` duplicate handling policy:

```
curl -k \  
-H "Authorization: Bearer ${CDE_TOKEN}" \  
-X POST "${CDE_JOBS_API}/admin/import" \  
-H "accept: application/json" \  
-H "Content-Type: multipart/form-data" \  
-F "file=@/path/to/cde-etl-jobs-backup.zip;type=application/zip" \  
-F duplicatehandling=rename \  
| jq
```

For Web UI

Before you begin

- You must have a valid backup file to restore from. For instructions on backing up CDE jobs, see [Backing up Cloudera Data Engineering jobs](#).

Steps

1. In the Cloudera Data Platform (CDP) management console, click the Data Engineering tile and click Overview.
2. In the CDE Services column, select the service containing the virtual cluster where you want to restore the jobs.
3. In the Virtual Clusters column on the right, click the View Jobs icon for the virtual cluster you want to restore to.
4. Click Jobs in the left menu.
5. Click the vertical ellipses menu at the top right, and then click Restore Jobs.
6. Click Choose a zip file.
7. Browse to the ZIP file containing the backup of jobs and resources you want to restore, and then click Open.
8. Click Select to restore the backup.

Result

The jobs and resources from the backup file are restored using the `rename` duplicate handling policy. If a job name in the backup conflicts with an existing job, the restore operation keeps the existing job and renames the restored job by appending a numerical identifier to the job name.