

Managing Jobs and Pipelines in Cloudera AI

Date published: 2020-07-16

Date modified: 2025-10-31

CLOUDERA

Legal Notice

© Cloudera Inc. 2026. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Creating a Job.....	4
Job Retry parameters.....	6
Creating a Pipeline.....	7
Viewing Job History.....	8
Legacy Jobs API (Deprecated).....	8

Creating a Job

This topic describes how to automate analytics workloads with a built-in job and pipeline scheduling system that supports real-time monitoring, job history, and email alerts.

A job automates the action of launching an engine, running a script, and tracking the results, all in one batch process. Jobs are created within the purview of a single project and can be configured to run on a recurring schedule. You can customize the engine environment for a job, set up email alerts for successful or failed job runs, and email the output of the job to yourself or a colleague.

Jobs are created within the scope of a project. When you create a job, you select a script to run as part of the job, and create a schedule for when the job is planned to run. Optionally, you can configure a job to be dependent on another existing job, thus creating a pipeline of tasks to be accomplished in a sequence.

Important:

- The script files and any other job dependencies must exist within the scope of the same project.
- For a recurring cron job to execute successfully, both the job creator's user account and the dedicated Service Account (Run As User) must remain active. For security and auditing reasons, if either the job creator's account or the Service Account is disabled or deleted, the recurring cron job will fail to run.

For Cloudera AI UI

1. Navigate to the project for which you want to create a job.
2. On the left-hand sidebar, click **Jobs**.
3. Click **New Job**.
4. Enter a Name for the job.
5. In **Run Job as**, if the job is to run in a service account, select **Service Account** and choose the account from the dropdown menu.
6. In **Script**, select a script to run for this job by clicking on the folder icon. You can select a script from a list of files that are already part of the project. To upload more files to the project, see *Managing Project Files*.
7. In **Arguments**, enter arguments to provide to the script.

Use the environment variable **JOB_ARGUMENTS** to access the arguments in a runtime-agnostic way.

For non-PBJ Runtimes, the content of this field is available as standard command line arguments.

8. Depending on the code you are running, select a Runtime Kernel for the job from the available options: Python 3.
9. Select one of the following options for a Schedule for the job runs:
 - **Manual** - Select this option if you plan to run the job manually each time.
 - **Recurring** - Select this option if you want the job to run in a recurring pattern every X minutes, or on an hourly, daily, weekly or monthly schedule. Set the recurrence interval with the drop-down buttons.

As an alternative, select the **Use a cron expression** checkbox and enter a Unix-style cron expression to set the interval. The expression must have five fields, specifying the minutes, hours, day of month, month, and day of week. If the cron expression is deselected, the default schedule indicated in the drop-down list takes effect.

- **Dependent** - Use this option when you are building a pipeline of jobs to run in a predefined sequence. From a dropdown list of existing jobs in this project, select the job that this one depends on. Once you have configured a dependency, this job will run only after the preceding job in the pipeline has completed a successful run.

10. Select a Resource Profile to specify the number of cores and memory available for each session.
11. Optionally, in the **Timeout in Minutes** field, enter a timeout value in minutes.

12. Configure the Job Retry Settings options at each job level.

Note: The Job Retry feature is available from Cloudera AI 1.5.5 SP1 or higher versions.

- a. Select Enable Retry to enable a retry run for the job.
- b. Configure the following Job Retry settings:

- Maximum Retry – The maximum number of retry attempts which can be triggered for a single job run in case of continuous failure of retry job runs.

The minimum value is 1.

- Retry Delay (minutes) – The delay between two consecutive retry job runs for a failed instance of the run.

The minimum value is 1 minute.

- Retry Conditions – Different options can be configured to control the terminal states of a job run that trigger a retry. The Retry process completes as soon as at least one (or more) option is selected.

Select at least one of the following criteria if Retry is enabled, but you can select any combination of the following Retry Conditions options:

- Script Failure – Runs the Retry process for user script failures if the user script exits with a non-zero exit code after the execution of the script.
- System Failure – Runs the Retry process for any kind of system- or engine-related failures not including user script failures.
- Timed-out Runs – Runs the Retry process for timed-out job runs.
- Skipped Runs – Runs the Retry process for skipped job runs.

13. Click Set environment variables if you want to set any values to override the overall project environment variables.**14.** Specify a list of Job Report Recipients to whom you can send email notifications with detailed job reports for job success, failure, or timeout. You can send these reports to yourself, your team if the project was created under a team account, or any other external email addresses.**15.** Add any Attachments such as the console log to the job reports that will be emailed.**16.** Click the Create Job button.

You can use the API v2 to schedule jobs from third party workflow tools. For more information, see *Using the Jobs API* as well as the *Cloudera AI API v2* tab.

For Cloudera AI API v2

Create a job using the API by following the settings included in the code:

```
job_body = cmlapi.CreateJobRequest()

# name and script
job_body.name = "my job name"
job_body.script = "pi.py"

# arguments
job_body.arguments = "arg1 arg2 \"all arg 3\""

# engine kernel
job_body.kernel = "python3" # or "r", or "scala"

# schedule
# manual by default
# for recurring/cron:
job_body.schedule = "* * * * 5" # or some valid cron string
# job retry settings
```

```
job_body.retry_enabled= true
job_body."max_retry"= 1
job_body."retry_delay"= 1
job_body.retry_for_script_failure= true
job_body.retry_for_skipped_runs= true
job_body.retry_for_system_failure= true
job_body.retry_for_timedout_runs= true

# for dependent (don't set both parent_job_id and schedule)
job_body.parent_job_id = "abcd-1234-abcd-1234"

# resource profile (cpu and memory can be floating point for partial)
job_body.cpu = 1 # one cpu vcore
job_body.memory = 1 # one GB memory
job_body.nvidia_gpu = 1 # one nvidia gpu, cannot do partial gpus

# timeout
job_body.timeout = 300 # this is in seconds

# environment
job_body.environment = {"MY_ENV_KEY": "MY_ENV_VAL", "MY_SECOND_ENV_KEY": "MY_SECOND_ENV_VAL"}

# attachment
job_body.attachments = ["report/1.txt", "report/2.txt"] # will attach /home/cds/w/report/1.txt and /home/cds/w/report/2.txt to emails

# After setting the parameters above, create the job:
client = cmlapi.default_client("host", "api key")
client.create_job(job_body, project_id="id of project to create job in")
```

**Note:**

The Job Retry feature and the Job retry settings are available from Cloudera AI 1.5.5 SP1 or higher releases.

For more examples of commands related to jobs, see *Using the Jobs API*.

Related Information

[Managing Project Files](#)

[Using the Jobs API](#)

[Legacy Jobs API \(Deprecated\)](#)

Job Retry parameters

The Job Retry feature is available from Cloudera AI 1.5.5 SP1 or higher releases. Configure Job Retry settings with the help of the Job Retry parameters.

If the Administrator configures the Job Retry settings, the specified values automatically populate the fields in the new job creation form when a user creates a job. In this case the Job Retry settings act as default values that the administrator can recommend to users.

If the Administrator does not configure the Job Retry settings, the fields remain blank in the new job creation form.

In both cases, users have the flexibility to customize the values during job creation or update them later through the job settings page.

Table 1: Job Retry parameters

Parameter	Description
Maximum Retry	<p>The Maximum Retry parameter defines the maximum number of retries that can be performed for a single failed job run. Retries continue until either a job run succeeds or the total number of retries reaches the maximum retry count specified by the user.</p> <p>Setting the Maximum Retry option to a high value can result in higher resource usage.</p> <p>The value must always be greater than 0.</p>
Retry Delay	<p>The Retry Delay value defines the time between each subsequent retry attempt in minutes.</p> <p>The Retry Delay period ensures that the transient errors, (for example, temporary network or resource outage, get fixed without overloading the system with job run requests.</p> <p>If you encounter transient issues, set the Retry Delay parameter to a higher value.</p> <p>If you address script failures and transient issues are not a concern, a lower value for the retry delay can be configured.</p> <p>If you have time-sensitive Jobs, set the Retry Delay parameter to a smaller value to trigger the retry at a faster pace.</p> <p>The value must always be greater than 0 and the minimum retry delay value is 1 minute.</p>
Retry Conditions	<p>The Retry Conditions parameter controls the terminal states of a job run that trigger a retry.</p> <p>Select minimum one of the criteria if Retry is enabled, but you can select any combination of the Retry Conditions options. The Retry process completes as soon as any of the selected criteria is met.</p> <p>The following Retry Conditions options can be enabled:</p> <ul style="list-style-type: none"> Script Failure – It runs the Retry process for user script failures if the user script exits with a non-zero exit code after the execution of the script. System Failure – It runs the Retry process for any kind of system- or engine-related failures not including user script failures. Timed-out Runs – It runs the Retry process for timed-out job runs. Skipped Runs – It runs the Retry process for skipped job runs. <p>The timeout value must be set to a reasonable duration. If the value is too short, each retry will encounter the same limit, potentially resulting in a continuous timeout # retry # timeout loop.</p>
Limit Concurrent Retries*	<p>*This parameter can only be set by the Administrator.</p> <p>It defines the limit value for the Maximum Concurrent Retry Limit parameter. The Maximum Concurrent Retry Limit parameter specifies the maximum number of job retry runs that can execute concurrently across the entire workbench, regardless of the total number of jobs running.</p> <p>If the maximum limit value defined as Maximum Concurrent Retry Limit is reached, any additional job retry runs are rescheduled until the number of active retry runs falls below the limit.</p> <p>Enable a hard limit only if job retry runs are consuming excessive resources, otherwise, avoid setting a hard limit.</p> <p>Administrators can set the Maximum Concurrent Retry Limit value if the Limit Concurrent Retries option is enabled.</p>

Creating a Pipeline

This topic describes how to create a scheduled pipeline of jobs within a project.

About this task

As data science projects mature beyond ad hoc scripts, you might want to break them up into multiple steps. For example, a project may include one or more data acquisition, data cleansing, and finally, data analytics steps. For such

projects, Cloudera AI allows you to schedule multiple jobs to run one after another in what is called a pipeline, where each job is dependent on the output of the one preceding it.

The Jobs overview presents a list of all existing jobs created for a project along with a dependency graph to display any pipelines you've created. Job dependencies do not need to be configured at the time of job creation. Pipelines can be created after the fact by modifying the jobs to establish dependencies between them. From the job overview, you can modify the settings of a job, access the history of all job runs, and view the session output for individual job runs.

Let us take an example of a project that has two jobs, Read Weblogs and Write Weblogs. Given that you must read the data before you can run analyses and write to it, the Write Weblogs job should only be triggered after the Read Weblogs job completes a successful run. To create such a two-step pipeline:

Procedure

1. Navigate to the project where the Read Weblogs and Write Weblogs jobs were created.
2. Click Jobs.
3. From the list of jobs, select *WRITE WEBLOGS*.
4. Click the Settings tab.
5. Click on the Schedule dropdown and select Dependent. Select Read Weblogs from the dropdown list of existing jobs in the project.
6. Click Update Job.

Viewing Job History

You can view the history for jobs that run within a project.

Procedure

1. Navigate to the project where the job was created.
The Project page displays.
2. Click Jobs on the left navigation page.
3. Select the relevant job.
4. Click the History tab. You see a list of all the job runs with some basic information, such as who created the job, run duration, and status. Click individual runs to see the session output for each run.
5. If Job Retry is enabled, click the Retry tag to view the failed run that triggered the retry.



Note:

The **Job Retry** feature is available from Cloudera AI 1.5.5 SP1 or higher versions.

On the Job History page, job runs that were retried are marked with a **Retry** tag. This option is visible only if retries are enabled and have been triggered.

Legacy Jobs API (Deprecated)

This topic demonstrates how to use the legacy API to launch jobs.

Cloudera AI exposes a legacy REST API that allows you to schedule jobs from third-party workflow tools. You must authenticate yourself before you can use the legacy API to submit a job run request. The Jobs API supports HTTP Basic Authentication, accepting the same users and credentials as Cloudera AI.



Note: The Jobs API is now deprecated. See *Cloudera AI API v2* and API v2 usage for the successor API.

Legacy API Key Authentication

Cloudera recommends using your legacy API key for requests instead of your actual username/password so as to avoid storing and sending your credentials in plaintext. The legacy API key is a randomly generated token that is unique to each user. It must be treated as highly sensitive information because it can be used to start jobs using the API. To look up your Cloudera AI legacy API key:

1. Sign in to Cloudera AI.
2. From the upper right drop-down menu, switch context to your personal account.
3. Click Settings.
4. Select the API Key tab.

The following example demonstrates how to construct an HTTP request using the standard **basic authentication** technique. Most tools and libraries, such as Curl and Python Requests, support basic authentication and can set the required Authorization header for you. For example, with curl you can pass the legacy API key to the --user flag and leave the password field blank.

```
curl -v -XPOST http://cdsw.example.com/api/v1/<PATH_TO_JOB> --user
  "<LEGACY_API_KEY>:"
```

To access the API using a library that does not provide Basic Authentication convenience methods, set the request's Authorization header to Basic `Basic <LEGACY_API_KEY_ENCODED_IN_BASE64>`. For example, if your API key is `uysgxtj7jzkps96njextnxmq05usp0b`, set Authorization to Basic `dXlzZ3h0ajdqemtwczk2bmpleHRueHhtcTA1dXNwMGI6`.

Starting a Job Run Using the API

Once a job has been created and configured through the Cloudera AI web application, you can start a run of the job through the legacy API. This will constitute sending a POST request to a job start URL of the form: `http://cdsw.example.com/api/v1/projects/<$USERNAME>/<$PROJECT_NAME>/jobs/<$JOB_ID>/start`.

To construct a request, use the following steps to derive the username, project name, and job ID from the job's URL in the web application.

1. Log in to the Cloudera AI web application.
2. Switch context to the team/personal account where the parent project lives.
3. Select the project from the list.
4. From the project's Overview, select the job you want to run. This will take you to the job Overview page. The URL for this page is of the form: `http://cdsw.example.com/<$USERNAME>/<$PROJECT_NAME>/jobs/<$JOB_ID>`.
5. Use the `$USERNAME`, `$PROJECT_NAME`, and `$JOB_ID` parameters from the job Overview URL to create the following job start URL: `http://cdsw.example.com/api/v1/projects/<$USERNAME>/<$PROJECT_NAME>/jobs/<$JOB_ID>/start`.

For example, if your job Overview page has the URL `http://cdsw.example.com/alice/sample-project/jobs/123`, then a sample POST request would be of the form:

```
curl -v -XPOST http://cdsw.example.com/api/v1/projects/alice/sample-project/jobs/123/start \
--user "<API_KEY>:" --header "Content-type: application/json"
```

Note that the request must have the Content-Type header set to application/json, even if the request body is empty.

Setting Environment Variables

You can set environment variables for a job run by passing parameters in the API request body in a JSON-encoded object with the following format.

```
{  
  "environment": {  
    "ENV_VARIABLE": "value 1",  
    "ANOTHER_ENV_VARIABLE": "value 2"  
  }  
}
```

The values set here will override the defaults set for the project and the job in the web application. This request body is optional and can be left blank.

Be aware of potential conflicts with existing defaults for environment variables that are crucial to your job, such as PATH and the Cloudera AI variables.

Sample Job Run

As an example, let's assume user Alice has created a project titled Risk Analysis. Under the Risk Analysis project, Alice has created a job with the ID, 208. Using curl, Alice can use her API Key (uysgxtj7jzkps96njextnxxmq05usp0b) to create an API request as follows:

```
curl -v -XPOST http://cdsw.example.com/api/v1/projects/alice/risk-analysis/jobs/208/start  
--user "uysgxtj7jzkps96njextnxxmq05usp0b:" --header "Content-type: application/json"  
--data '{ "environment": { "START_DATE": "2017-01-01", "END_DATE": "2017-01-31" } }'
```

In this example, START_DATE and END_DATE are environment variables that are passed as parameters to the API request in a JSON object.

In the resulting HTTP request, curl automatically encodes the Authorization request header in base64 format.

```
* Connected to cdsw.example.com (10.0.0.3) port 80 (#0)  
* Server auth using Basic with user 'uysgxtj7jzkps96njextnxxmq05usp0b'  
> POST /api/v1/projects/alice/risk-analysis/jobs/21/start HTTP/1.1  
> Host: cdsw.example.com  
> Authorization: Basic dxLzz3h0ajdqemtwczk2bmpleHRueHhtcTA1dXNwMGI6  
> User-Agent: curl/7.51.0  
> Accept: */*  
> Content-type: application/json  
>  
< HTTP/1.1 200 OK  
< Access-Control-Allow-Origin: *  
< Content-Type: application/json; charset=utf-8  
< Date: Mon, 10 Jul 2017 12:00:00 GMT  
< Vary: Accept-Encoding  
< Transfer-Encoding: chunked  
<  
{  
  "engine_id": "cwg6wclmg0x482u0"  
}
```

You can confirm that the job was started by going to the Cloudera AI web application.

Starting a Job Run Using Python

To start a job run using Python, Cloudera recommends using [Requests](#), an HTTP library for Python; it comes with a convenient API that makes it easy to submit job run requests to Cloudera AI. Extending the Risk Analysis example

from the previous section, the following sample Python code creates an HTTP request to run the job with the job ID, 208.

Python 2

```
# example.py

import requests
import json

HOST = "http://cdsw.example.com"
USERNAME = "alice"
API_KEY = "uysgxtj7jzkps96njextnxxmq05usp0b"
PROJECT_NAME = "risk-analysis"
JOB_ID = "208"

url = "/".join([HOST, "api/v1/projects", USERNAME, PROJECT_NAME, "jobs",
JOB_ID, "start"])
job_params = {"START_DATE": "2017-01-01", "END_DATE": "2017-01-31"}
res = requests.post(
    url,
    headers = {"Content-Type": "application/json"},
    auth = (API_KEY, ""),
    data = json.dumps({"environment": job_params})
)

print "URL", url
print "HTTP status code", res.status_code
print "Engine ID", res.json().get('engine_id')
```

When you run the code, you should see output of the form:

```
python example.py
```

```
URL http://cdsw.example.com/api/v1/projects/alice/risk-analysis/jobs/208/sta
rt
HTTP status code 200
Engine ID r11w5q3q589ryg9o
```

Limitations

- Cloudera AI does not support changing your legacy API key, or having multiple API keys.
- Currently, you cannot create a job, stop a job, or get the status of a job using the Jobs API.

Related Information

[API v2 usage](#)

[Basic Access Authentication](#)

[Creating a Pipeline](#)

[Environment Variables](#)

[Cloudera AI API v2](#)