

Data Engineering deployment architecture

Date published: 2020-07-30

Date modified: 2024-11-12



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

Recommendations for scaling Cloudera Data Engineering deployments.....	4
Apache Airflow scaling and tuning considerations.....	5
Guidelines for Virtual Cluster upkeep.....	6
Guidelines for database upkeep.....	6

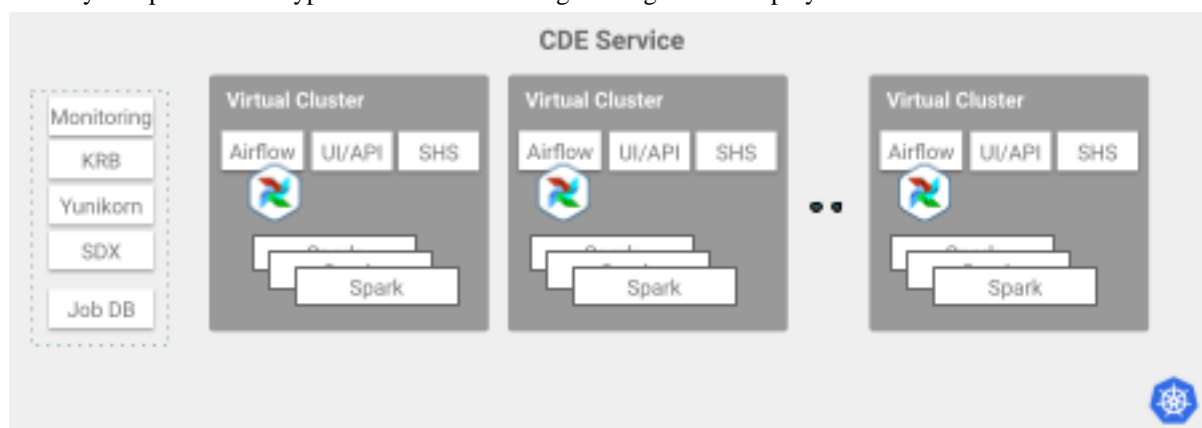
Recommendations for scaling Cloudera Data Engineering deployments

Your business might experience a sudden increase or drop in demand due to which your Cloudera Data Engineering deployment needs to autoscale. You can scale your Cloudera Data Engineering deployment by either adding new instances of a Cloudera Data Engineering service or Virtual Cluster, or by adding additional resources to the existing ones.

You can scale your Cloudera Data Engineering deployment:

- Vertically - More resources are provisioned within the same instance of a Cloudera Data Engineering service or Virtual Cluster.
- Horizontally - New instances of Cloudera Data Engineering service or Virtual Cluster are provisioned.

The key components of a typical Cloudera Data Engineering service deployment are:



Virtual Clusters provide an isolated autoscaling compute capacity to run Spark and Airflow jobs. You can use Virtual Clusters to isolate individual teams or lines of business by using user-based Access Control Lists (ACLs).

Guidelines for scaling Virtual Clusters

- Each Virtual Cluster requires infrastructure capacity to run various services such as Airflow, API server, and Spark-History-Server (SHS).

Recommendation: Do not scale horizontally beyond 50 Virtual Clusters within the same Cloudera Data Engineering service.

- Virtual Clusters can actively run hundreds of parallel jobs. In certain scenarios, it might be required to simultaneously submit multiple jobs as per the schedule or due to a burst in demand. In these scenarios, the API server enforces guardrails and limits the number of simultaneous Spark job submissions to 60. Once the jobs move from submission to running state, you can submit more jobs.



Note: The Airflow jobs are not impacted by the submission limit of 60 simultaneous jobs.

Recommendation: Distribute simultaneous submission of jobs over time or horizontally scale across multiple Virtual Clusters.

Job submission rate guardrails

When jobs are submitted to the Cloudera Data Engineering API server of a particular Virtual Cluster, it takes time and resources, known as preparing and starting states, before they begin running. This process is called the job submission intake process. To ensure proper resourcing and handling of incoming jobs, guardrails have been set up. By default, the guardrail, or limit is set to 60 simultaneous job submissions. Simultaneous incoming Spark job

submissions that exceed 60 return a 429 error message to the client. The example error message is: Failed to submit. Too many requests.

Recommendation:

- Incorporate CDE Clients error handling for this error. The CDE CLI exit code is 77. The CDE CLI receives an HTTP 429 response code for the request to the Runtime API Server's REST API. For more information on exit codes, see Cloudera Data Engineering CLI exit codes.
- If needed, increase or decrease the guardrail limit during the Virtual Cluster creation by setting the "limiting.simultaneousJobSubmissions" configuration.

Related Information

[Creating virtual clusters](#)

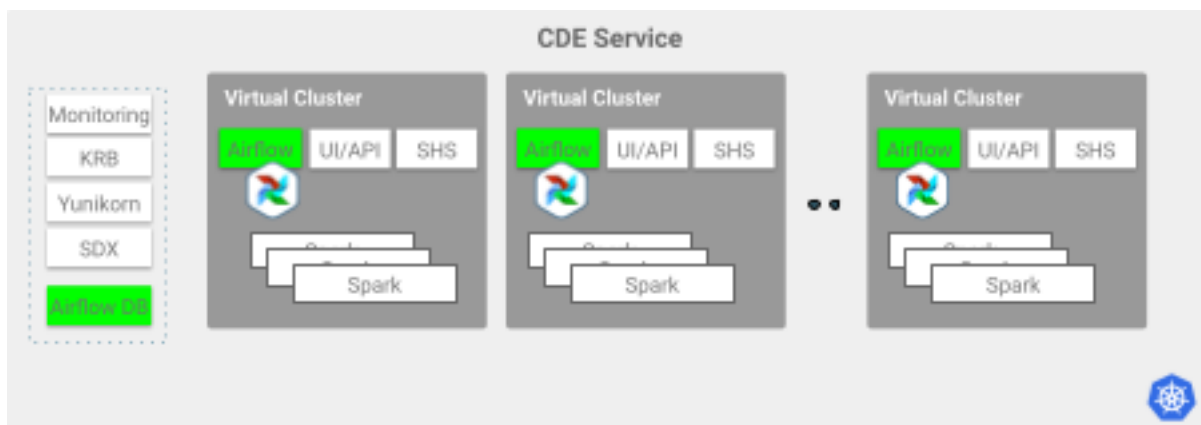
[Cloudera Data Engineering CLI exit codes](#)

Apache Airflow scaling and tuning considerations

There are certain limitations related to the deployment architecture, and guidelines for scaling and tuning of the deployment, that you must consider while creating or running Airflow jobs (DAGs).

CDE deployment architecture

If you use Airflow to schedule and/or develop multi-step pipelines, you must consider the following deployment limitations to scale your environment.



These guidelines describe when to scale horizontally, given total loaded DAGs and/or number of parallel (concurrent) tasks. These apply to all cloud providers.

CDE service guidelines

The number of Airflow jobs that can run at the same time is limited by the number of parallel tasks that are triggered by the associated DAGs. The number of tasks can reach up to 250-300 when running in parallel per CDE service.

Recommendation: Create more CDE services to increase Airflow tasks concurrency beyond the limits noted above.

Virtual Cluster guidelines

- Each Virtual Cluster has a maximum parallel task execution limit of 250. This applies to both Amazon Web Services (AWS) and Azure.



Note: It is recommended that you do not exceed 300 parallel tasks across all Virtual Clusters.

- For AWS: The recommended maximum number of Airflow jobs loaded per virtual cluster is 1500. As the number of jobs increases, the job page loading time experiences latency along with Airflow job creation time.
- For Azure: The recommended maximum number of Airflow jobs loaded per virtual cluster is 300 to 500. Higher capacity can be achieved but will require manual configuration changes.

Recommendation: Create more Virtual Clusters to load additional Airflow DAGs.

- The suggested number of Airflow jobs created in parallel is three or less. For jobs that are created in parallel, specify a DAG and/or a task_group for each Airflow operator.

Recommendation: Distribute the creation of Airflow jobs so that you do not create more than three jobs at the same time.

Guidelines for Virtual Cluster upkeep

There are upkeep guidelines for Cloudera Data Engineering (CDE) Spark History Server (SHS) that you'll need to consider.

Lifecycle configuration of CDE Spark event logs

The number of Spark event logs (spark.eventLog.enabled), that are produced by Spark jobs that run via CDE Virtual Cluster, grows indefinitely with each new CDE Spark run. These event logs are not automatically deleted and are stored on the object store under <CDP env storage location>/dex/<Service ID>/<VC ID>/eventlog/.

Some examples of the event log location can look like the following:

- For Amazon Web Services (AWS): s3a://dex-storage-bucket/datalake/logs/dex/cluster-2xv14pfp/rdw8q2sh/eventlog/
- For Azure: abfs://logs@dexstorageaccount.dfs.core.windows.net/dex/cluster-4p54mk8j/22bnm99g/eventlog/

Because the number of event logs continuously increases, the time from when the CDE job finishes and when the Spark UI is available for this run on the Virtual Cluster UI may increase. The delay is most apparent in Virtual Clusters with 6,000 or more completed Spark job runs.

To avoid delays in event log availability after CDE job runs, you can configure an object store lifecycle policy so that event logs are deleted automatically on the object store. For more information about an Amazon S3 lifecycle policy, see [Setting lifecycle configuration on a bucket](#) linked below. For more information about Azure lifecycle management policies, see [Configure a lifecycle management policy](#) linked below.



Note: The lifecycle policy rule prefix does not include a bucket or storage account name.

Related Information

[Setting lifecycle configuration on a bucket](#)

[Configure a lifecycle management policy](#)

Guidelines for database upkeep

There are database upkeep guidelines for Cloudera Data Engineering (CDE) that you'll need to consider.



Note: Despite the end of service for MySQL v8.0.28 in June 2023, you can continue using CDE with MySQL v8.0.28. Although MySQL will continue to work with your databases after end of service, troubleshooting support will not be available. Therefore, it is recommended that you upgrade to a supported version of MySQL.

Upgrade MySQL for AWS

Upgrade MySQL for AWS to ensure you are using a supported version to work with CDE.



Note: If you are using an in-place upgrade, your MySQL version will also be updated.

1. In the AWS console, navigate to RDS > cluster-id > Modify.
2. Select the version that you want to upgrade from the Database engine version drop-down list.
3. Click Continue.
4. Click Apply Immediately to apply changes.

Increasing database size on Azure

Azure databases can fill up with logs and entries with every job that runs in CDE Virtual Clusters. When the database fills up, increase the size of your database.

1. Ensure that there are no jobs running.
2. Go to the Azure Portal.
3. Locate the Azure SQL database named <cluster id>, for example, cluster-2cmmds8q.
4. Navigate to the Pricing Tier section.
5. Move the slider to the desired value.



Note: The slider cannot be reduced.

6. Click OK and wait for the scaling to complete.
7. Resume running your jobs.

Increasing database size on Amazon Web Services

Amazon Web Services (AWS) databases can fill up with logs and entries with every job that runs in CDE Virtual Clusters. When the database fills up, increase the size of your database.

1. Go to the AWS Console.
2. Navigate to the RDS Service page.
3. Click Databases and use the filter to find your cluster id, for instance, cluster-w8d65npx.
4. Select the target database.
5. In the Database homepage in the top right-hand corner, click Modify.
6. Scroll down to the database size settings and set the Allocated Storage Property to the desired value.
7. Click Continue.
8. Set the desired timeframe for maintenance.
9. Click Continue and wait for the database status to become available.
10. Resume running your jobs.