

## Managing virtual clusters

Date published: 2020-07-30

Date modified: 2025-11-20





# Legal Notice

© Cloudera Inc. 2025. 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

<b>Creating virtual clusters.....</b>	<b>4</b>
<b>Viewing and managing virtual cluster details.....</b>	<b>7</b>
Applying user and group access for virtual clusters.....	8
Editing user and group access for virtual clusters.....	10
Cloudera Data Engineering Virtual Cluster Access Controls Troubleshooting.....	11
Managing Virtual Cluster-level Spark configurations (Technical Preview).....	11
<b>Editing virtual clusters.....</b>	<b>13</b>
<b>Locking and unlocking a Cloudera Data Engineering Service.....</b>	<b>16</b>
<b>Overview of suspending and resuming Cloudera Data Engineering virtual clusters [Technical Preview].....</b>	<b>16</b>
<b>Deleting virtual clusters.....</b>	<b>17</b>
<b>Accessing non-public Cloudera Data Engineering virtual clusters via SOCKS proxy.....</b>	<b>18</b>



## Creating virtual clusters

In Cloudera Data Engineering, a virtual cluster is an individual auto-scaling cluster with defined CPU and memory ranges. Jobs are associated with virtual clusters, and virtual clusters are associated with an environment. You can create as many virtual clusters as you need. See Recommendations for scaling Cloudera Data Engineering deployments linked below.

### Before you begin



**Important:** The user interface for Cloudera Data Engineering 1.17 and above has been updated. The left-hand menu was updated to provide easy access to commonly used pages. The steps below will vary slightly Cloudera Data Engineering version prior to 1.17, for example, the Overview page has been replaced with the Home page. The new Home page still displays Virtual Clusters, but now includes quick-access links located at the top for the following categories: Jobs, Resources, and Download & Docs.



**Note:** Virtual Cluster installation fails when Ozone S3 gateway proxy is enabled. Ozone S3 gateway proxy gets enabled when more than one Ozone S3 Gateway is configured in the Cloudera Base on premises cluster.

As a workaround, add the `127.0.0.1 s3proxy-<environment-name>.<private-cloud-control-plane-name>-services.svc.cluster.local` entry in the `/etc/hosts` file of all nodes in the Cloudera Base on premises cluster where the Ozone S3 gateway is installed. For example, if the private cloud environment name is `cdp-env-1` and private cloud control plane name is `cdp`, then add the following entry in the `/etc/hosts` file:

```
127.0.0.1 s3proxy-cdp-env-1.cdp-services.svc.cluster.local
```

To create a virtual cluster, you must have an environment with Cloudera Data Engineering enabled.

### Procedure

1. In the Cloudera console, click the Data Engineering tile. The Cloudera Data Engineering Home page displays.
2. Click Administration in the left navigation menu, select the environment you want to create a virtual cluster in.



3.

In the Virtual Clusters column, click  at the top right to create a new virtual cluster.

If the environment has no virtual clusters associated with it, the page displays a Create DE Cluster button that launches the same wizard.

### Cluster Details

Cluster Name \*

Service\*

What do you intend to use this cluster for? ⓘ

☒ **Core (Tier 1)**  
Operational deployment & monitoring of batch based transformations

☐ **All Purpose (Tier 2)** NEW  
Develop using interactive sessions and deploy both batch & streaming (coming soon) workloads.

Unsure which to select? [Review our feature comparison](#) or [pricing breakdown](#) ↗

Select Spark Version

Auto-Scale Max Capacity

CPU

Memory (GB)

---

### Additional Options

Enable Iceberg tables, Restrict user access

4. Enter a Cluster Name.

Cluster names must include the following:

- Begin with a letter
- Be between 3 and 30 characters (inclusive)
- Contain only alphanumeric characters and hyphens

5. Select the Service to create the virtual cluster in.

The environment you selected before launching the wizard is selected by default, but you can use the wizard to create a virtual cluster in a different Cloudera Data Engineering service.



6. Select one of the following Cloudera Data Engineering cluster types:



**Note:** These tiers are available in Cloudera Data Engineering 1.19 and above only.

- Core (Tier 1)
- All Purpose (Tier 2)



**Important:** For both Microsoft Azure and Amazon Web Services, to avoid being charged for All purpose compute nodes, even if you do not create an All purpose Virtual Cluster:

Navigate to the Service Details page > Configurations tab > Capacity & Costs > Autoscale Range and set the minimum value of All purpose On-demand Instances to 0.

For more information, see [Cloudera Data Engineering cluster types](#).

7. Select the Spark Version to use in the virtual cluster. You cannot use Spark 2 and Spark 3 in the same virtual cluster, but you can have separate Spark 2 and Spark 3 virtual clusters within the same Cloudera Data Engineering service. While you can have virtual clusters with different Spark 3.x versions, a single virtual cluster can only use one Spark version.

On the Cloudera Data Engineering UI, you can view the supported Apache Airflow runtime component version and Spark runtime component version on the VC creation and details pages.

The following screenshot illustrates creating a new Virtual Cluster, where you can select the required runtime component versions from a drop-down menu on the UI.

Administration / Create a Virtual Cluster

### Virtual Cluster Details

Setup your virtual cluster by completing the following fields. These settings cannot be changed once your cluster has been created.

\* Name

\* Service

\* Spark Version  
 Python 3.11 | Iceberg 1.5.2 | Java 17 | Scala 2.12.18

Need more information on compatibility of runtime components? [Review our compatibility matrix](#)

Airflow Runtime Version  
 Python 3.11



**Important:** If you switch from the Redhat UBI image to the security hardened image, note that the updated component versions and libraries can impact the Cloudera Data Engineering operation. Cloudera strongly recommends testing the new implementation in a non-production environment first.

For backward compatibility, you can continue using the Redhat UBI image. The security hardened image is the recommended image, however, note that it requires testing and validation.

For more information, see [Security hardened Spark image migration guide](#).

For more information, see [Compatibility for Cloudera Data Engineering and Runtime components](#).

8. Use the Auto-Scale Max Capacity sliders to set the maximum number of CPU cores and the maximum memory in gigabytes. The cluster will scale up and down as needed to run the submitted Spark applications.
9. Optional for spot instances enabled at the Cloudera Data Engineering service level: From the Driver and Executors will run on drop-down menu, select whether you want to run drivers and executors on spot instances or on-demand instances. By default, the driver runs on on-demand instances, and the executors run on spot instances. For SLA-bound workloads, select On-demand. For non-SLA workloads, Cloudera recommends leaving the



default configuration to take advantage of the cost savings afforded by spot instances. For more information, see [Cloudera Data Engineering Spot Instances](#).

10. Select Restrict Access to add access control for the virtual cluster. You can search for users to add by name or email address. You can manage users using the Cloudera Management Console. For more information, see [Managing user access and authorization](#).
11. Click Configure Email Alerting (Technical Preview) if you want to receive notification mails. The email configuration options appear.
  - a) You must provide at least Sender Email Address and SMTP Host information.
  - b) Test SMTP Configs: Click Test SMTP Configs to test the configurations set for SMTP. This helps you to test the SMTP configuration before creating the cluster.



**Note:** You will have the option to enable email alerts when you create or edit a job after the Virtual Cluster is created, but step 13 must be completed first.

12. Click Create.

### Results

On the Cloudera Data Engineering Home page, select the environment to view the virtual cluster initialization status. You can also click the three-dot menu for the virtual cluster to view the logs.

### Related Information

[Recommendations for scaling Cloudera Data Engineering deployments](#)

## Viewing and managing virtual cluster details

You can view details such as configuration, metrics, and logs for existing virtual clusters. You can also view jobs associated with a cluster, clone a cluster, and delete a cluster



**Important:** The user interface for Cloudera Data Engineering 1.17 and above has been updated. The left-hand menu was updated to provide easy access to commonly used pages. The steps below will vary slightly, for example, the Overview page has been replaced with the Home page. The new home page still displays Virtual Clusters, but now includes quick-access links located at the top for the following categories: Jobs, Resources, and Download & Docs.

To view cluster details:

1. In the Cloudera console, click the Data Engineering tile. The Cloudera Data Engineering Home page displays.
2. Click Administration in the left navigation menu. The Administration page displays.
3. In the Services column, select the environment containing the virtual cluster you want to manage.
4. In the Virtual Clusters column on the right, click the Cluster Details icon for the virtual cluster you want to manage.
5. On the Configuration tab, you can view details about the cluster, including the cluster name, Cloudera environment, and the CPU and memory auto-scale ranges.

The virtual cluster Overview page provides details about the virtual cluster, and includes links to several resources, such as the API documentation, and the Jobs API URL. You can switch between the following three tabs to view additional information:

### Configuration

The Configuration tab lists details about the cluster, including the cluster name, Cloudera environment, and the CPU and memory auto-scale ranges. You can modify the virtual cluster quotas (CPU/memory) dynamically. Cloudera Data Engineering provides an Edit option to make the configuration changes dynamically which may take a few minutes to update. For more information, see [Editing virtual clusters](#) linked below.

### Charts



The Charts tab displays charts of virtual cluster metrics, such as network and filesystem usage.



**Note:** If your browser blocks third-party cookies, Charts may not display correctly. If this happens, add \*.cloudera.com as an exception in your browser settings.

### Security

The Security tab displays the Access Management that you can set for the virtual cluster. See Applying user and group access for virtual clusters for more details.

### Logs

The Logs tab displays the latest log entries for the virtual cluster.

### Related Information

[Editing virtual clusters](#)

## Applying user and group access for virtual clusters

Learn how to restrict or grant access to a virtual cluster for specific users or groups that you specify in Cloudera Data Engineering.

### Procedure

1. In the Cloudera console, click the Data Engineering tile. The Cloudera Data Engineering Home page displays.
2. Click Administration in the left navigation menu. The Administration page displays.
3. In the Services column, select the environment containing the virtual cluster you want to manage.
4. In the Virtual Clusters column on the right, click the Cluster Details icon for the virtual cluster you want to update security for.



- Click the Security Access Add User or Group .

## DE\_Americas

Administration / Virtual Cluster

STATUS	VERSION	CLUSTER-ID	TYPE	CREATED BY
Maintenance	1.1	[redacted]	All Purpose	[redacted]

DATA LAKE ENVIRONMENT

go [redacted] vs

Configuration Monitoring **Security** New Airflow New

Access **New**

Access Management

Grant individual users or groups  
To Create a user or group you w

**Add User or Group**

The Add User(s) or Group(s) screen is displayed.

- In the Search for User(s) or Group(s) field, enter the name or email of the user(s) or group name(s).
- Under Set Role(s) drop-down, select Read-only Admin or Admin.

The Admin option can access endpoints including all jobs, resources, and sessions. The Read-only Admin option can only view all cluster activities but cannot modify the cluster or activities.

- Click Add.

The Access Management screen is displayed with the users and groups that you added. Only one role can be set each time, but selection can be modified by editing once the user or groups are added.

- To add more users or groups, click Add.

The Add User(s) or Group(s) screen is displayed.

- Repeat steps 6 through 8.

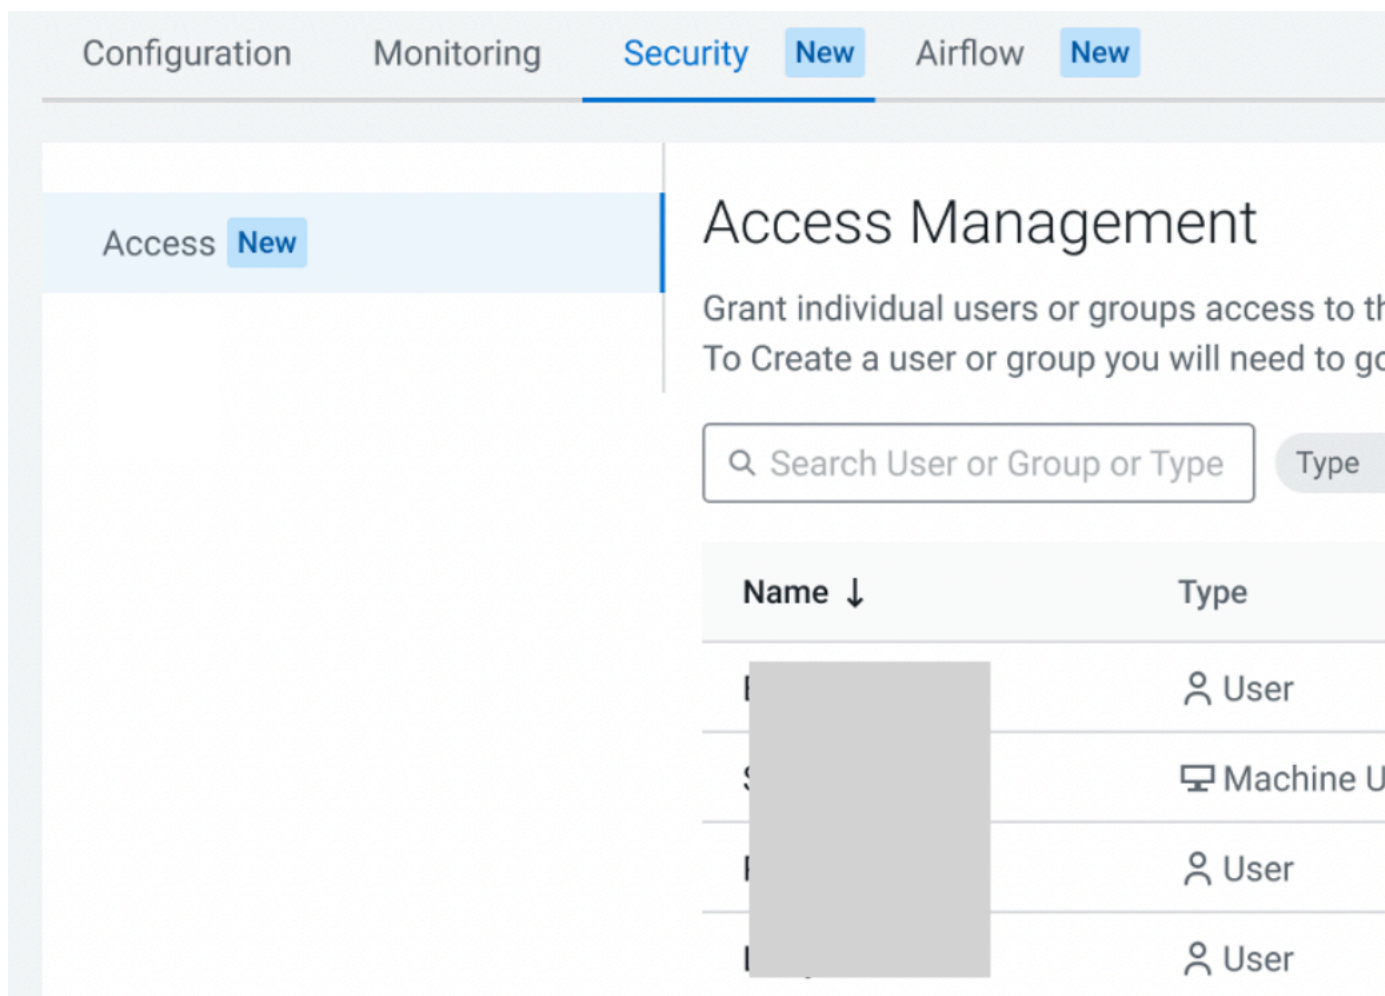


## Editing user and group access for virtual clusters

Learn how to edit or delete a user or group access that you have added for a virtual cluster in Cloudera Data Engineering.

### Procedure

1. In the Cloudera console, click the Data Engineering tile. The Cloudera Data Engineering Home page displays.
2. Click Administration in the left navigation menu. The Administration page displays.
3. In the Services column, select the environment containing the virtual cluster you want to manage.
4. In the Virtual Clusters column on the right, click the Cluster Details icon for the virtual cluster you want to update security for.
5. Click the Security Access .



6. Locate the name of the user or group that you want to edit.
7. Click the Pencil icon to edit the current user or group.  
The Add User(s) or Group(s) screen is displayed.
8. If you want to update the user or group, in the Search for User(s) or Group(s) field, enter the name or email of the user(s) or group name(s).
9. If you want to edit the role, under Set Role(s) drop-down, select Read-only Admin or Admin.  
The Admin option can access endpoints including all jobs, resources, and sessions. The Read-only Admin option can only view all cluster activities but cannot modify the cluster or activities.



10. If you want to delete the user or group, locate the name or group of the user that you want to delete, and click the Trash icon.

## Cloudera Data Engineering Virtual Cluster Access Controls Troubleshooting

Learn about troubleshooting scenarios for user and group access issues in Cloudera Data Engineering Virtual Clusters (VCs).

### A user is added to a group with access to a VC, but the user still cannot access the VC

As the user is added to the group, do a user synchronization. For information about synchronizing users in an environment, see [Performing user sync](#).



**Note:** If the user still cannot access the VC, it might be due to internal caching. Wait for five minutes and try again.

### A user or group is not displayed in the drop-down list while adding roles

If the user or group exists in the Cloudera account, but the user is not displayed in the drop-down list, it is very likely that the user is given a role already. Check the current access list.

### Adding access controls while creating a VC

From the UI, you can add access controls only after the VC is created. Using CLI or API, you can add access controls during VC creation.

### A user is added, but if the user tries to access the VC, the UI shows a 403 error

The user needs to have either a DEUser or a DEAdmin role in the environment irrespective of the roles in the VC. For more information, see [Giving access to Data Engineering users](#).

### User access works, but group access does not

If the group access does not work as expected, do a user synchronization in the environment and wait for five minutes. If the user access still does not work, a miscommunication between the authorization server and the LDAP server can be the cause. To find such errors, access the kubernetes cluster and check for logs of the dex-authz pod in the dex-base- namespace.

## Managing Virtual Cluster-level Spark configurations (Technical Preview)

Learn how to add, edit, or delete the Spark configurations at a Virtual Cluster (VC)-level so that the configurations apply to all the jobs in the Virtual Cluster, by default. You can edit the Spark configurations using the Cloudera Data Engineering UI, the CDP CLI, or the Cloudera Data Engineering API.

### About this task

Earlier, the Spark configurations were added to the jobs at the job-level. If you had to apply the common configurations for all the jobs, then you had to add the configurations for each job individually.

Using this feature, you can apply the Spark configurations in the Virtual Cluster (VC) to all the jobs by default.

### Before you begin

For information on how to manage Virtual Cluster-level Spark configurations using the API, see [Managing Virtual Cluster-level Spark configurations using the API](#).

### Procedure

Edit the Spark configurations using the Cloudera Data Engineering UI, the CLI, or the Cloudera Data Engineering API.

**For Cloudera Data Engineering UI**



- a. Sign in to Cloudera Data Engineering UI and click Administration.
- b. Select the relevant Cloudera Data Engineering Service.
- c. Select the relevant Virtual Cluster in the Cloudera Data Engineering Service and click the Cluster Details option.
- d. Click Configurations Spark .

A Spark page appears.

- e. Under Configurations, Enter Spark Configurations as key-value pairs.

For example

```
is.config.vc=true
is.config.job=false
```

- f. Click Apply Changes.

Result: The Spark configurations are updated and applied to the VC and the jobs in the Virtual Cluster. You can validate them by running the jobs and then verifying the Spark configurations in the Configuration tab in that particular job runtime.

### For CDP CLI

- To create a VC with the Spark configurations, run the following command:

```
cdp de create-vc \
  --name [***VIRTUAL-CLUSTER-NAME***] \
  --cluster-id [***CLUSTER-ID***] \
  --cpu-requests [NUMBER-OF-CPUS-REQUIRED***] \
  --memory-requests [NUMBER-OF-MEMORY UNITS-REQUIRED***] \
  --spark-configs \ [***KEY1***]=[***VALUE1***],[***KEY2***]=[***VALUE2***]
```

For example

```
cdp de create-vc \
  --name test-vc \
  --cluster-id cluster-sq4rpnr2 \
  --cpu-requests 20 \
  --memory-requests 20 \
  --spark-configs \
  spark.executor.instances=2,spark.dynamicAllocation.maxExecutors=15,key1=vall
```

- To update the Spark configurations of the VC, run the following command:

```
cdp de update-vc \
  --cluster-id [***CLUSTER-ID***] \
  --vc-id [***VC-ID***] \
  --spark-configs \ [***KEY1***]=[***VALUE1***][***KEY2***]=[***VALUE2***]
```

For example

```
cdp de update-vc \
  --cluster-id cluster-sq4rpnr2 \
  --vc-id dex-app-w2rhk89r \
  --spark-configs \ spark.executor.instances=2,spark.dynamicAllocation.maxExecutors=15,key1=vall
```

- To discard or delete all the Spark configurations of the VC, run the following command:

```
cdp de update-vc \
```



```
--cluster-id      [ ***CLUSTER-ID*** ] \
--vc-id           [ ***VC-ID*** ] \
--discard-spark-configs
```

For example

```
cdp de update-vc \
  --cluster-id      cluster-sq4rpnr2 \
  --vc-id           dex-app-w2rhk89r \
  --discard-spark-configs
```

## Editing virtual clusters

You can edit virtual cluster configurations in the UI or use a script to apply configurations to multiple jobs at once.



**Important:** The user interface for Cloudera Data Engineering 1.17 and above has been updated. The left-hand menu was updated to provide easy access to commonly used pages. The steps below will vary slightly, for example, the Overview page has been replaced with the Home page. The new home page still displays Virtual Clusters, but now includes quick-access links located at the top for the following categories: Jobs, Resources, and Download & Docs.

### For Edit in UI

To edit virtual cluster details:

1. In the Cloudera console, click the Data Engineering tile. The Cloudera Data Engineering Home page displays.
2. Click Administration in the left navigation menu. The Administration page displays.
3. In the Services column, select the environment containing the virtual cluster you want to manage.
4. In the Virtual Clusters column on the right, click the Cluster Details icon for the virtual cluster you want to manage.
5. On the Configuration tab, you can modify the virtual cluster quotas (CPU/memory) dynamically. Cloudera Data Engineering provides an Edit option to make the configuration changes dynamically which may take a few minutes to update.

### For Edit with Script

As an Administrator, you can further customize the virtual cluster by applying additional configurations all at once using the `cde-utils` script. Click [cdp-cde-utils.sh](#) to download the script. Customizations include applying spark configurations to all jobs, enabling and disabling gang scheduling, and initiating faster autoscaling times for nodes in an autoscaler cluster for Amazon Web Services (AWS). See the following `cde-utils` script and options to apply these changes:



**Important:** Before using the script, the user must export the kubeconfig file to ensure that the virtual cluster is accessible via `kubectl`.

### Apply Spark configurations to all new jobs

This script allows users to add and update default configurations for all Spark jobs all at once. The options required are as follows:

- `-h <host_name>`: You must provide this option, and paste the domain of the URL into a text editor to identify the endpoint host. For example, if the URL is `https://xyz.cde-vjgd8ksl.dsp-azur.xcu2-8y8x.dev.cldr.work/dex/ui/jobs`, then paste the following :

```
xyz.cde-vjgd8ksl.dsp-azur.xcu2-8y8x.dev.cldr.work
```



- `-c <spark_default_configurations>`: You must provide option `-c <spark_configuration>` to add all the spark-configurations inside the single quotes ( ' ') with a delimiter comma (,) at the virtual cluster level. There must be at least one configuration in the command. For example:

```
./cde-utils.sh add-spark-config-in-virtual-cluster -h <host_name> -c 'spark.driver.supervise="false",spark.executor.instances="2"'
```

The ,user can add multiple configurations as shown above in the example: ' "key1"="val1", "key2"="val2", ... '. The , <host\_name> is is determined by the provided information in the host name option.

### Enabling and disabling gang scheduling

This script enables and disables gang scheduling at the virtual cluster level for all jobs that are inside the virtual cluster. The options required are as follows:

- `-h <host_name>`: You must provide this option, and paste the domain of the URL into a text editor to identify the endpoint host. For example, if the URL is `https://xyz.cde-vjgd8ksl.dsp-azur.xcu2-8y8x.dev.cldr.work/dex/ui/jobs`, then paste the following:

```
xyz.cde-vjgd8ksl.dsp-azur.xcu2-8y8x.dev.cldr.work
```

- `--gang-scheduling <gang_scheduling_status>`: You must provide this option to toggle the status of the gang scheduling at the virtual cluster level. You can only use enable or disable options. To disable gang scheduling, see the following example:

```
./cde-utils.sh add-spark-config-in-virtual-cluster -h <host_name> --gang-scheduling disable
```

The , <host\_name> is determined by the provided information in the host name option.



**Important:** You must provide either configurations option (`-c`) or gang-scheduling option (`--gang-scheduling`) for this sub-command. You may also provide both options.

### Initiating faster autoscaling times

You can increase the speed of autoscaling times for nodes in an autoscaler cluster for Amazon Web Services (AWS) by using a script. This script allows you to speed up the scale down progress for situations where you scale up to hundreds of nodes. This is a cost-saving measure. The options required for scaling down nodes are as follows:

- `-i <interactive_menu>`: You can use this option to use an interactive menu that uses the arrow keys to navigate through the menu options. For example:

```
./cde-utils.sh edit-cluster-autoscaler -i
```



**Important:** In the edit-cluster-autoscaler sub command, you can provide all the other options except `-i <interactive_menu>`, or only provide `-i <interactive_menu>` to change the value of other options using the interactive menu.

- `--scale-down-delay-after-add`: Reduce the value of this option which specifies the interval after scaling up when scale down evaluation resumes. This allows you to speed up the process. The default value is 10 minutes. Permitted values are 10 seconds, 5 minutes, and 1 hour. The format is as follows: `<time>[s / m / h]`. For Example:

```
./cde-utils.sh edit-cluster-autoscaler --scale-down-delay-after-add 10m
```

- `--scale-down-delay-after-delete`: You can decrease the value for this option which is the duration of time after node deletion when the scale down evaluation resumes. The default is 10 seconds.



Permitted values are 10 seconds, 5 minutes, and 1 hour. The format is as follows: <time>[s / m / h]. For example:

```
./cde-utils.sh edit-cluster-autoscaler --scale-down-delay-after-delete 10s
```

- `--scale-down-delay-after-failure`: You can decrease the value for this option which is a duration of time after a scale down failure when the scale down evaluation resumes. The default value is 3 minutes. Permitted values are 10 seconds, 5 minutes, and 1 hour. The format is as follows: <time>[s / m / h]. For example:

```
./cde-utils.sh edit-cluster-autoscaler --scale-down-delay-after-failure 3m
```

- `--scale-down-unneeded-time`: To speed up the process, the user can decrease the value of this option which specifies the amount of time before a node qualifies for scale down. The default value is 10 minutes. Permitted values are 10 seconds, 5 minutes, and 1 hour. The format is as follows: <time>[s / m / h]. For example:

```
./cde-utils.sh edit-cluster-autoscaler --scale-down-unneeded-time 10m
```

- `--unremovable-node-recheck-timeout`: You can reduce the value for this option which is the duration of the timeout before the unremovable node is checked again. The default value is 5 minutes. Permitted values are 10 seconds, 5 minutes, and 1 hour. The format is as follows: <time>[s / m / h]. For example:

```
./cde-utils.sh edit-cluster-autoscaler --unremovable-node-recheck-timeout 5m
```

### Example use cases with commands

These example use cases and commands show how you can use the `cde-utils` script.

- Initiate faster autoscaling times with the command below. The order of options is not significant in this command:

```
./cde-utils.sh edit-cluster-autoscaler --scale-down-delay-after-add 1m --scale-down-delay-after-failure 1m --scale-down-delay-after-delete 5s --scale-down-unneeded-time 1m --unremovable-node-recheck-timeout 1m
```

- Add a single Spark configuration at the virtual cluster level:

```
./cde-utils.sh add-spark-config-in-virtual-cluster -h <host_name> -c spark.executor.instances="2"
```

- Add multiple spark configurations at the virtual cluster level:

```
./cde-utils.sh add-spark-config-in-virtual-cluster -h <host_name> -c 'spark.driver.supervise="false",spark.executor.instances="2"'
```

- Enable gang scheduling:

```
./cde-utils.sh add-spark-config-in-virtual-cluster -h <host_name> --gang-scheduling enable
```

### Related Information

[Accessing the Kubernetes dashboard](#)



[Support Video: Enabling kubectrl for Cloudera Data Engineering](#)

## Locking and unlocking a Cloudera Data Engineering Service

Locking a Cloudera Data Engineering Service freezes the configuration of a Service and its corresponding Virtual Clusters. When a Service is locked, you are unable to edit, add, or delete a Service and its Virtual Clusters.


### About this task

You may want to lock a Service in scenarios where you don't want configuration changes to occur during a planned upgrade of Cloudera Data Engineering. The lock and unlock options will display interchangeably, depending on the state that the Service is in. For example, if the Service is unlocked, you will only see the Lock Service option when you complete the steps below. Once your Service is locked, you will see the option to Unlock it as you complete the steps below.

### Before you begin

Note that locking a Service is a preview feature starting in Cloudera Data Engineering 1.19. Ensure that you are on Cloudera Data Engineering 1.19 and above.

### Procedure

1. In the Cloudera console, click the Data Engineering tile. The Cloudera Data Engineering Home page displays.
2. Click Administration.
3. Locate the Service that you want to lock or unlock.
4. Click  on the Service, and select the Lock Service or Unlock Service option.
5. Enter the name of the Service that you want to lock or unlock as this confirms your choice.
6. Click the appropriate Lock Service or Unlock button.

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

The virtual cluster-level suspend and resume feature enables you to temporarily pause one or more Cloudera Data Engineering Virtual Clusters (VCs) during idle periods.



**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.

### Benefits of suspending virtual clusters

Suspending a VC scales down all associated infrastructure and job pods to zero. Depending on the distribution of the pods, suspending a VC can release one or more nodes, leading to cost savings. Cloudera has observed that suspending a larger number of VCs offers significant cost savings, primarily on the compute resources.

During suspension, all VC configurations, as well as jobs, and job runs are preserved, similar to hibernating a laptop during idle periods. While a VC is in suspended state, it does not allow the submission or the execution of jobs on the respective VC. However, other VCs remain operational.

To reduce operational costs, you can suspend a larger number of VCs during weekends, or overnight on weekdays.

The suspend and the resume operation can be initiated by:



- A user
- A scheduled event
- An automatic action



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

### Prerequisites for suspending and resuming virtual clusters

To suspend or resume VCs, you need to:

- Have DEAdmin role
- Request the DE\_ENABLE\_VC\_SUSPEND\_RESUME entitlement

### Limitations

- You can suspend and resume VCs through the CDP CLI and the Cloudera Data Engineering API, but not through the Cloudera Data Engineering UI.
- Often you might not save on infrastructure nodes when you suspend a VC. This heavily depends on the distribution of infrastructure pods across nodes.
- When you suspend a VC, the following operations are not supported:
  - Upgrading the Cloudera Data Engineering service
  - Backing up the Cloudera Data Engineering service
  - Backing up the suspended VC jobs (you can back up jobs from other running VCs)
- Attempting to suspend multiple VCs simultaneously can sometimes lead to failures. As a best practice, suspend VCs in a sequential manner.



**Note:** You do not need to resume VCs in a sequential manner, you can resume multiple VCs in parallel.

### Related Information

[Suspending and resuming Cloudera Data Engineering virtual clusters using the API \[Technical Preview\]](#)

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

## Deleting virtual clusters

If you no longer need a particular virtual cluster, you can delete it. Deleting a virtual cluster instantly deletes all metadata associated with the cluster, including job configurations and log files. Do not delete a virtual cluster unless you are certain you no longer need access to any of the applications, jobs, or log files. Job run history is maintained separately, and is not deleted even if you delete a virtual cluster.

### Before you begin



**Important:** The user interface for Cloudera Data Engineering 1.17 and above has been updated. The left-hand menu was updated to provide easy access to commonly used pages. The steps below will vary slightly, for example, the Overview page has been replaced with the Home page. The new home page still displays Virtual Clusters, but now includes quick-access links located at the top for the following categories: Jobs, Resources, and Download & Docs.

### Procedure

1. On the Cloudera Data Engineering Overview page, select the environment containing the virtual cluster that you want to delete.



2. In the Virtual Clusters column for the selected environment, click the menu icon at the top right of the virtual cluster that you want to delete, and then click Delete.



**Warning:** Deleting a virtual cluster deletes all metadata associated with the cluster, including job configurations and log files. Do not do this unless you are certain that you no longer need any of these.

3. Confirm that you want to delete the cluster by entering the cluster name and then clicking Delete.

### Results

The virtual cluster will show the status as deleted and will be fully removed from the Cloudera Data Engineering UI in about a week. You can create a new virtual cluster with the same name as the deleted one.

## Accessing non-public Cloudera Data Engineering virtual clusters via SOCKS proxy

A SOCKS proxy server, or jump host, allows your web browser to connect securely to a non-public Cloudera Data Engineering endpoint. When you connect to the Cloudera Data Engineering endpoint using its internal IP, the browser connects to the proxy server, which performs the required SSH tunneling, and forwards the request to the internal Cloudera Data Engineering endpoint.

### About this task

This topic describes how to configure a SOCKS proxy to access Cloudera Data Engineering Virtual Clusters on non-publicly routable VPCs.



**Note:** Using the below approach all the web traffic in the launched browser instance will be routed via the jump host.

### Procedure

1. In a VPC with a public subnet, which can also route to the target VPC where the Cloudera Data Engineering Load Balancer is running, create an EC2 instance for your SOCKS server (for example, my-ec2-jump-host) with a public IP and an SSH key-pair (for example, my-key-file.pem).

The instance does not need to be large, for example on AWS a micro instance is usually sufficient, but may need to be larger depending on how many users are using it concurrently.

Depending on whether you want multiple users to share the jump host or have everyone create their own server, pick the SSH key pair for the instance accordingly.



**Note:** Make sure to apply a locked down security group to your jump host instance to ensure only users from trusted IP ranges can SSH into the server.

Related AWS Documentation: [Getting Started with Amazon EC2 Instances](#), [Amazon EC2 Key Pairs](#)

- a) Set up a SOCKS proxy server via SSH to the EC2 instance using below command:

```
nohup ssh -i "SSH-KEY-PAIR" -CND <SOCKS-PROXY-PORT> ec2-user@<JUMP-HOST-PUBLIC-IP> &
```



```
Example: nohup ssh -i "my-key-file.pem" -CND 8157 ec2-user@<public ip for my-ec2-jump-host> &
```

where

- nohup(optional) is a POSIX command to ignore the HUP (hangup) signal so that the proxy process is not terminated automatically if the terminal process is later terminated.
- my-key-file.pem is the private key you used to create the EC2 instance where the SOCKS server is running.
- C sets up compression.
- N suppresses any command execution once established.
- D 8157 sets up the SOCKS 5 proxy on the port. The port number 8157 in this example is arbitrary, but must match the port number you specify in your browser configuration in step 2.
- ec2-user is the AMI username for the EC2 instance. The AMI username can be found in the details for the instance displayed in the AWS Management Console on the Instances page under the Usage Instructions tab.
- <public ip for my-ec2-jump-host> is the public IP address of the EC2 instance running the SOCKS server you created above.
- & (optional) causes the SSH connection to run as an operating system background process, independent of the command shell. Without the &, you leave your terminal open while the proxy server is running and use another terminal window to issue other commands.

## 2. Configure Your Browser to Use the Proxy.

Using Google Chrome

By default, Google Chrome uses system-wide proxy settings on a per-profile basis. To get around that you can start Chrome using the Command Line and specify the following:

- The SOCKS proxy port to use (must be the same value used in step 1, 8157 in example above)
  - The profile to use (new profile created in the example below)
- a) Run the following commands to create a new profile and launch a new instance of Chrome that does not interfere with any currently running instance.

For Linux:

```
/usr/bin/google-chrome \
--user-data-dir="$HOME/chrome-with-proxy" \
--proxy-server="socks5://localhost:<SOCKS-PROXY-PORT>"
```

For MacOS

```
"/Applications/Google Chrome.app/Contents/MacOS/Google Chrome" \
--user-data-dir="$HOME/chrome-with-proxy" \
--proxy-server="socks5://localhost:<SOCKS-PROXY-PORT>"
```

For Windows

```
"C:\Program Files (x86)\Google\Chrome\Application\chrome.exe" ^
--user-data-dir="%USERPROFILE%\chrome-with-proxy" ^
--proxy-server="socks5://localhost:<SOCKS-PROXY-PORT>"
```

## Results

These steps enable you to navigate to any Cloudera Data Engineering virtual cluster in the browser launched using SOCKS proxy.