

Managing the Embedded Container Service (ECS)

Date published: 2023-12-16

Date modified: 2024-10-18



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

The Embedded Container Service (ECS).....	4
Configuring the Embedded Container Service.....	4
Upgrading RHEL to a new minor version.....	4
Adding hosts to a Embedded Container Service Cluster.....	5
Starting, stopping, restarting, and refreshing Embedded Container Service Clusters.....	22
Starting a Embedded Container Service Cluster.....	23
Stopping a CDP Private Cloud Data Services Cluster.....	23
Restarting a Embedded Container Service Cluster.....	23
Rolling Restart of an Embedded Container Service Cluster.....	23
Configuring Restart for an Embedded Container Service cluster.....	24
Refreshing a Embedded Container Service Cluster.....	24
Monitoring Embedded Container Service Clusters.....	24
Viewing Health Status.....	24
Viewing the Kubernetes Dashboard.....	25
Viewing the Private Cloud Management Console.....	25
Performing maintenance of a single host in the Embedded Container Service cluster.....	25
Performing the maintenance of all hosts in the Embedded Container Service cluster.....	26
Configuring a containerized cluster with SELinux.....	26
Decommissioning ECS Hosts.....	27
Dedicating ECS nodes for specific workloads.....	28
Specifying racks for ECS clusters.....	29
ECS unified time zone.....	42
Adjusting the expiration time of ECS cluster certificates.....	43
Configuring multiple Base clusters with one ECS cluster.....	49
GPU node labeling on ECS.....	51

The Embedded Container Service (ECS)

Cloudera Manager provides tools for managing and monitoring the CDP Private Cloud Embedded Container Service.

The Embedded Container Service (ECS) service enables you to run CDP Private Cloud Data Services by creating container-based clusters in your data center. In addition to the option to use OpenShift, which requires that you deploy and manage the Kubernetes infrastructure, you can also deploy a Embedded Container Service cluster, which creates and manages an embedded Kubernetes infrastructure for use with CDP Private Cloud Data Services. Installing, configuring, and managing OpenShift is not required. You only need to provide hosts on which to install the service and Cloudera Manager sets up the Embedded Container Service cluster and also provides management and monitoring of the cluster.

When you create an Embedded Container Service cluster, two new services are added to the cluster:

- Embedded Container Service (ECS) service. The ECS service has two roles:
 - ECS Server -- runs on a single host in the Embedded Container Service cluster.
 - ECS Agent -- runs on all hosts except the host running the Server role in the Embedded Container Service Cluster.
- Docker service. The Docker service has a single role:
 - Docker Server -- runs on all hosts in the Embedded Container Service Cluster.

Configuring the Embedded Container Service

You use Cloudera Manager to configure the Embedded Container Service and clusters.

Procedure

1. Open the Cloudera Manager Admin Console
2. From the Home page, Click on the Embedded Container Service Cluster
3. Click the Hosts, ECS service, or the Docker service links.
4. Click the Configuration tab.
5. Use the Filters or Search functions to locate the configuration property you are looking for.
6. Enter your change.
7. Click Save Changes.

Related Information

[Modifying Configuration Properties Using Cloudera Manager](#)

Upgrading the RHEL Operating System to a new minor version

After installing CDP Private Cloud Data Services on a particular RHEL OS, you can now upgrade RHEL OS to a new minor version. For example, in a RHEL 8.x OS series, you can upgrade from RHEL 8.6 to RHEL 8.8 new minor version.

About this task

You must perform this task on all Embedded Container Service (ECS) hosts.

Before you begin

Verify the RHEL OS version by performing the following steps:

1. Log in to Cloudera Manager as an Administrator, then navigate to the ECS cluster.

2. Go to the **Home Status** tab.
3. Click the Actions menu to the right of the Embedded Container Service (ECS) cluster name and select **Stop**.
4. Click the **Stop** option in the confirmation screen.

The **Command Details** window shows the progress of the services.

5. SSH into a ECS cluster host as a root user.
6. Verify the OS version by running the following command:

```
cat /etc/redhat-release
```

Procedure

1. Upgrade RHEL from the installed version to the desired minor version on the ECS host. Use the operating system upgrade procedures provided by your RedHat operating system vendor to download and upgrade RHEL.
For example, you can upgrade from RHEL 8.6 version to RHEL 8.8 minor version.
2. Verify the upgraded OS version by running the following command:

```
cat /etc/redhat-release
```

3. Log in to Cloudera Manager as an Administrator.
4. Go to the **Home Status** tab.
5. Click the Actions menu to the right of the Embedded Container Service cluster name and select **Start**.
6. Click the **Start** option.

The **Command Details** window shows the progress of the services.

Wait for all the pods to start. The wait time depends on the number of nodes in the cluster.

Adding hosts to a Embedded Container Service Cluster

You can add hosts to a Embedded Container Service (ECS) cluster to increase capacity and performance.

About this task

Procedure

1. On the Cloudera Manager home page, click the ECS Cluster, then select Actions > Add Hosts.

The screenshot shows the Cloudera Manager interface for the ECS cluster 152-b883. The left sidebar contains navigation links for Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, and Data Services. The main content area shows the cluster status, including a list of hosts (3 Hosts), DOCKER, and ECS. The 'Actions' dropdown menu is open, showing options like Start, Stop, Restart, Refresh Cluster, Upgrade Cluster, Inspect Hosts in Cluster, Rolling Restart, Rename Cluster, Enter Maintenance Mode, and View Maintenance Mode Status. The 'Add Hosts' option is highlighted. Below the status section, there are three charts: Cluster CPU, Cluster Disk IO, and Cluster Network IO, each showing performance metrics over time.

2. On the Add Hosts page, click Add Hosts to Cluster and select the ECS Cluster, then click Continue.

The screenshot shows the 'Add Hosts' wizard in Cloudera Manager. The wizard explains that it allows installing the Cloudera Manager Agent on new hosts. There are two options: 'Add hosts to Cloudera Manager' and 'Add hosts to Cluster'. The 'Add hosts to Cluster' option is selected, and a dropdown menu shows the cluster '152-b883'. At the bottom right, there are 'Back' and 'Continue' buttons.

3. On the Specify Hosts page, hosts that have already been added to Cloudera Manager are listed on the Currently Managed Hosts tab. You can select one or more of these hosts to add to the ECS cluster.

Cloud
ERA

Manager

Parcels

Running Commands

Support

admin

7.11.3

1 Specify Hosts

2 Install Parcels

3 Inspect Hosts

4 Select Host Template

5 Deploy Client Config

Specify Hosts

Currently Managed Hosts (1/4 Selected)

New Hosts

These hosts do not belong to any clusters. Select some to form your cluster.

<input type="checkbox"/>	Hostname (FQDN) ↑	IP Address	Rack	Version	Cores
<input type="checkbox"/>	dh-centos79m-1.vpc.cloudera.com	10.65.202.225	/default	None	8
<input type="checkbox"/>	dh-centos79m-2.vpc.cloudera.com	10.65.203.223	/default	None	8
<input type="checkbox"/>	dh-centos79m-3.vpc.cloudera.com	10.65.202.91	/default	None	8
<input checked="" type="checkbox"/>	ecst-2.vpc.cloudera.com	10.65.203.79	/default	None	8


1 - 4 of 4

Cancel

← Back

Continue →

You can also click the New Hosts tab to specify one or more hosts that have not been added to Cloudera Manager. Enter a Fully Qualified Domain Name in the Hostname box, then click Search.

 **Note:** Click the pattern link under the Hostname box to display more information about allowed FQDN patterns.

Cloud
ERA

Manager

Parcels

Running Commands

Support

admin

7.11.3

1 Specify Hosts

2 Select Repository

3 Select JDK

4 Enter Login Credentials

5 Install Agents

6 Install Parcels

7 Inspect Hosts

8 Select Host Template

9 Deploy Client Config

Specify Hosts

Currently Managed Hosts (1/4 Selected)

New Hosts (1 Selected)

Hosts should be specified using the same hostname (FQDN) that they will identify themselves with.

Hostname

ecst-[1-2].vpc.cloudera.com

Hint: Search for hostnames or IP addresses using [pattern](#)

SSH Port

22

Search

2 hosts scanned, 2 running SSH.

<input type="checkbox"/>	Expanded Query	Hostname (FQDN) ↑	IP Address	Currently Managed	Result
<input checked="" type="checkbox"/>	ecst-1.vpc.cloudera.com	ecst-1.vpc.cloudera.com	10.65.196.65	No	Host was successfully scanned.
<input type="checkbox"/>	ecst-2.vpc.cloudera.com	ecst-2.vpc.cloudera.com	10.65.203.79	Yes	Host was successfully scanned.

1 - 2 of 2

Cancel

← Back

Continue →

After you have finished specifying the ECS hosts, click Continue.

7

4. On the Select Repository page, the applicable Cloudera Manager Agent repository location is selected by default. Click Continue.

CLUSTER CDEP Deployment from 2023-Oct-23 11:55

Add Hosts

- Specify Hosts
- Select Repository**
- Select JDK
- Enter Login Credentials
- Install Agents
- Install Parcels
- Inspect Hosts
- Select Host Template
- Deploy Client Config

Select Repository

Cloudera Manager Agent 7.11.3 (#46431848) needs to be installed on all new hosts.

Repository Location ☐ Cloudera Repository (Requires direct Internet access on all hosts.) ☒ Custom Repository

Example: `http://LOCAL_SERVER/cloudera-repos/cm7/7.11.3`

Do not include operating system-specific paths in the URL. The path will be automatically derived.

[Learn more at How to set up a custom repository.](#)

[Cancel](#) [Back](#) [Continue](#)

5. Select a JDK option on the Select JDK page, then click Continue.

Add Hosts

- Specify Hosts
- Select Repository
- Select JDK**
- Enter Login Credentials
- Install Agents
- Install Parcels
- Inspect Hosts
- Select Host Template
- Deploy Client Config

Select JDK

CDH Version	Supported JDK Version
7.1.9 and above	OpenJDK 8, 11, 17 or Oracle JDK 8, 11, 17
7.1.1 to 7.1.8	OpenJDK 8, 11 or Oracle JDK 8, 11
7.0 and above	OpenJDK 8 or Oracle JDK 8
6.3 and above	OpenJDK 8 or Oracle JDK 8
6.2	OpenJDK 8 or Oracle JDK 8
6.1 or 6.0	Oracle JDK 8
5.16 and above	OpenJDK 8 or Oracle JDK 8
5.7 to 5.15	Oracle JDK 8

1 - 8 of 8
[More details on supported JDK version.](#)

If you plan to use JDK 11 with CDH 7.1.x and above or JDK 17 with CDH 7.1.9 and above, you will need to install it manually on all hosts and then select the **Manually manage JDK** option below.

☐ Manually manage JDK

☒ Install a Cloudera-provided version of OpenJDK

By proceeding, Cloudera will install a supported version of OpenJDK version 8.

☐ Install a system-provided version of OpenJDK

By proceeding, Cloudera will install the default version of OpenJDK version 8 provided by the Operating System.

[Cancel](#) [Back](#) [Continue](#)

- On the Enter Login Credentials page, All hosts accept the same password is selected by default. Enter the user name in the SSH Username box, and type in and confirm the password. You can also select the All hosts accept the same private key option and provide the Private Key and passphrase.

The screenshot shows the Cloudera Manager interface during the 'Add Hosts' process. On the left is a dark sidebar with the Cloudera Manager logo and a navigation menu including 'Parcels', 'Running Commands', 'Support', and 'admin'. Below the menu is the version '7.11.3' and a double-left arrow icon. The main area is titled 'Add Hosts' with a timestamp 'CDEP Deployment from 2023-Oct-23 11:55'. A progress bar on the left lists steps: 'Specify Hosts', 'Select Repository', 'Select JDK', 'Enter Login Credentials' (current step), 'Install Agents', 'Install Parcels', 'Inspect Hosts', 'Select Host Template', and 'Deploy Client Config'. The 'Enter Login Credentials' section contains a warning: 'Root access to your hosts is required to install the Cloudera packages. This installer will connect to your hosts via SSH and log in either directly as root or as another user with password-less sudo/pbrun privileges to become root.' Below this are fields for 'SSH Username' (set to 'root'), 'Authentication Method' (radio buttons for 'All hosts accept same password' (selected) and 'All hosts accept same private key'), 'Password' (masked with dots), 'Confirm Password' (masked with dots), 'SSH Port' (set to '22'), and 'Simultaneous Installations' (set to '10' with a note: '(Running a large number of installations at once can consume large amounts of network bandwidth and other system resources)'). At the bottom are 'Cancel', 'Back', and 'Continue' buttons.

CLUSTER CLOUDERA Manager

7.11.3

Parcels
Running Commands
Support
admin

Add Hosts

CDEP Deployment from 2023-Oct-23 11:55

- Specify Hosts
- Select Repository
- Select JDK
- Enter Login Credentials**
- Install Agents
- Install Parcels
- Inspect Hosts
- Select Host Template
- Deploy Client Config

Enter Login Credentials

Root access to your hosts is required to install the Cloudera packages. This installer will connect to your hosts via SSH and log in either directly as root or as another user with password-less sudo/pbrun privileges to become root.

SSH Username

Authentication Method ☒ All hosts accept same password
☐ All hosts accept same private key

Password

Confirm Password

SSH Port


Simultaneous Installations
(Running a large number of installations at once can consume large amounts of network bandwidth and other system resources)

Cancel Back Continue

7. The Cloudera Manager agents are installed, and then the Install Parcels page appears. The selected parcel is downloaded to the Cloudera Manager server host, distributed, unpacked, and activated on the ECS cluster hosts. Click Continue.

The screenshot shows the Cloudera Manager interface during the 'Add Hosts' process. On the left is a dark sidebar with the Cloudera Manager logo and a list of navigation items: Parcels, Running Commands, Support, and an 'admin' user profile. The main content area is titled 'Add Hosts' and features a vertical progress bar on the left with steps: Specify Hosts, Select Repository, Select JDK, Enter Login Credentials, Install Agents, **Install Parcels** (current step), Inspect Hosts, Select Host Template, and Deploy Client Config. The 'Install Parcels' section displays the message: 'The selected parcels are being downloaded and installed on all the hosts in the cluster.' Below this, a progress bar for 'Embedded Container Service ...' shows four stages: 'Downloaded: 100%', 'Distributed: ...', 'Unpacked: 4/4', and 'Activated: 4/4'. At the bottom of the interface are 'Cancel', '← Back', and 'Continue →' buttons. A status bar at the very bottom shows the version '7.11.3' and a double-left arrow icon.

8. Review the Validations list on the Inspect Hosts page. If issues are detected, you can fix the issues, then click Run Again to repeat the host inspection. Click Continue.



Parcels

Running Commands

Support

admin

7.11.3

Add Hosts

Specify Hosts

Select Repository

Select JDK

Enter Login Credentials

Install Agents

Install Parcels

Inspect Hosts

Select Host Template

Deploy Client Config

CDEP Deployment from 2023-Oct-23 11:55

Inspect Hosts

[Run Again](#)

Status	Description
✓	Inspector ran on all 4 hosts.
✓	Individual hosts resolved their own hostnames correctly.
✓	No errors were found while looking for conflicting init scripts.
✓	No errors were found while checking /etc/hosts.
✓	All hosts resolved localhost to 127.0.0.1.
✓	All hosts checked resolved each other's hostnames correctly and in a timely manner.
✓	Host clocks are approximately in sync (within ten minutes).
✓	Host time zones are consistent across the cluster.
✓	No users or groups are missing.
✓	No conflicts detected between packages and parcels.
✓	No kernel versions that are known to be bad are running.
✓	No problems were found with /proc/sys/vm/swappiness on any of the hosts.
⚠	Transparent Huge Page Compaction is enabled and can cause significant performance problems. Run "echo never > /sys/kernel/mm/transparent_hugepage/defrag" and "echo never > /sys/kernel/mm/transparent_hugepage/enabled" to disable this, and then add the same command to an init script such as /etc/rc.local so it will be set on system reboot. The following hosts are affected: > View Details
⚠	Hue Python version dependency is satisfied. Starting with CDH 6, PostgreSQL-backed Hue requires Psycopg2 version to be at least 2.5.4, see the documentation for more information. The following hosts are missing a compatible version of the Psycopg2 library: > View Details
✓	A compatible version of the operating system is installed on the hosts in a Private Cloud Containerized Cluster.
✓	Ports 80 and 443 are available for use on the hosts in a Private Cloud Containerized Cluster.

[Cancel](#)

← Back
Continue →

9. The Select Host Template page lists available host templates. Click Create.

**Note:**

The following three steps describe how to create a host template to assign the Docker Server and Ecs Agent role groups to the new host. You can also select None and add these role instances after adding the new host to the cluster, as described at the end of this topic.

The screenshot shows the Cloudera Manager interface during the 'Add Hosts' process. On the left is a dark sidebar with the Cloudera Manager logo and a list of navigation items: Parcels, Running Commands, Support, and a user profile for 'admin'. The main area is titled 'Add Hosts' and features a progress bar with nine steps: Specify Hosts, Select Repository, Select JDK, Enter Login Credentials, Install Agents, Install Parcels, Inspect Hosts, **Select Host Template** (the current step), and Deploy Client Config. The 'Select Host Template' panel contains the text 'Select a host template to apply to the new hosts in order to populate them with role instances.' and a radio button labeled 'None' which is selected. Below the radio button is a 'Create...' button. At the bottom of the wizard are three buttons: 'Cancel', '< Back', and 'Continue >'. A status bar at the very bottom shows the version '7.11.3' and a double-left arrow icon. A small blue banner at the top right of the main area reads 'CDEP Deployment from 2023-Oct-23 11:55'.

10. On the Create New Host Template pop-up, enter a template name and select the Docker Server and Ecs Agent role groups, then click Create.

The screenshot shows the 'Add Hosts' page in Cloudera Manager. A modal window titled 'Create New Host Template For 152-b883' is open. The 'Template Name' field contains 'ecsworker'. Under 'Select Role Groups to Include:', the 'DOCKER' section has 'Docker Server' checked with a dropdown set to 'Docker Server Default Group'. The 'ECS' section has 'Ecs Agent' checked with a dropdown set to 'Ecs Agent Default Group', and 'Ecs Server' is unchecked. The modal has 'Cancel' and 'Create' buttons at the bottom right. The background shows the 'Add Hosts' progress bar and navigation buttons.

11. On the Select Host Template page, select the new template, then click Continue.

The screenshot shows the 'Add Hosts' page in Cloudera Manager, specifically the 'Select Host Template' step. The progress bar on the left indicates the current step is '8 Select Host Template'. The main content area is titled 'Select Host Template' and contains the instruction 'Select a host template to apply to the new hosts in order to populate them with role instances.' There are two radio buttons: 'None' and 'ecsworker', with 'ecsworker' selected. A 'Create...' button is next to the 'ecsworker' option. Below this, there is a checked checkbox for 'Start newly created roles after applying the host template.' At the bottom, there are 'Cancel', 'Back', and 'Continue' buttons.

12. The Apply Host Template page appears. After the roles have successfully started, click Continue.

CLOUDERA
Manager

Parcels

Running Commands

Support

admin

7.11.3

<<

Add Hosts

Specify Hosts

Select Repository

Select JDK

Enter Login Credentials

Install Agents

Install Parcels

Inspect Hosts

Select Host Template

Apply Host Template

Deploy Client Config

CDEP Deployment from 2023-Oct-23 11:55

Apply Host Template

Start Roles on Hosts When Free Command

Status Finished Dec 12, 10:20:41 PM 48.4s

Successfully started all the roles on selected hosts.

Completed 3 of 3 step(s).

Show All Steps

Show Only Failed Steps

Show Only Running Steps

> Wait for Service Commands	DOCKER	Dec 12, 10:20:41 PM	99ms
> Wait for Service Commands	ECS	Dec 12, 10:20:41 PM	100ms
> Starts all the roles on the selected hosts.		Dec 12, 10:20:41 PM	48.25s

Cancel

← Back

Continue →

13. The Deploy Client Config page appears. After all client configurations have been successfully deployed, click Finish.

CloudEra
Manager

Parcels

Running Commands

Support

admin

7.11.3

Specify Hosts

Select Repository

Select JDK

Enter Login Credentials

Install Agents

Install Parcels

Inspect Hosts

Select Host Template

Apply Host Template

10 Deploy Client Config

Add Hosts

Deploy Client Config

Deploy Client Configuration Command

Status Finished Context [152-b883](#) Dec 12, 10:26:12 PM 59ms

Successfully deployed all client configurations.

Completed 1 of 1 step(s).

Show All Steps Show Only Failed Steps Show Only Running Steps

Execute DeployClusterClientConfig for {} in parallel.

Dec 12, 10:26:12 PM

57ms

Cancel

Back Finish

14. The new host is listed on the ECS cluster Hosts page.

CloudEra
Manager

Search

Clusters

Hosts

Diagnostics

Audits

Charts

Replication

Administration

Data Services New

152-b883

Hosts

Configuration Add Hosts Review Upgrade Status Inspect Hosts in Cluster Inspect Cluster Network Performance

Search

Filters

Last Updated: Dec 12, 10:29:36 PM UTC

Filters

STATUS

Good Health 4

CLUSTERS

CORES

COMMISSION STATE

LAST HEARTBEAT

LOAD (1 MINUTE)

LOAD (5 MINUTES)

LOAD (15 MINUTES)

MAINTENANCE MODE

UPGRADE DOMAIN

RACK

SERVICE

Actions for Selected

Status	Name	IP	Roles	Tags	Commission State	Last He
<input type="checkbox"/>	<input checked="" type="checkbox"/>	dh-centos79-1.vpc.cloudera.com	10.65.203.160	2 Roles	Commissioned	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	dh-centos79-2.vpc.cloudera.com	10.65.194.119	2 Roles	Commissioned	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	dh-centos79-3.vpc.cloudera.com	10.65.194.114	2 Roles	Commissioned	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	ecst-1.vpc.cloudera.com	10.65.217.129	2 Roles 1 Tag	Commissioned	

1 - 4 of 4

- 15.** If your ECS hosts are running the CentOS 8.4, OEL 8.4, RHEL 7.9, or RHEL 8 operating systems, you must install iptables on all the ECS hosts.

For CentOS 8.4, OEL 8.4, or RHEL 8, run the following command on each ECS host:

```
yum --setopt=tsflags=noscripts install -y iptables
```

For RHEL 7.9, run the following command on each ECS host:

```
yum install -y iptables
```

16. If you did not apply a host template to assign roles, perform the following steps to assign the Docker Server and Ecs Agent role groups to the new host.

To assign the Docker Server role group:

- a. Click DOCKER on the ECS cluster home page, select Instances, then click Add Role Instances.

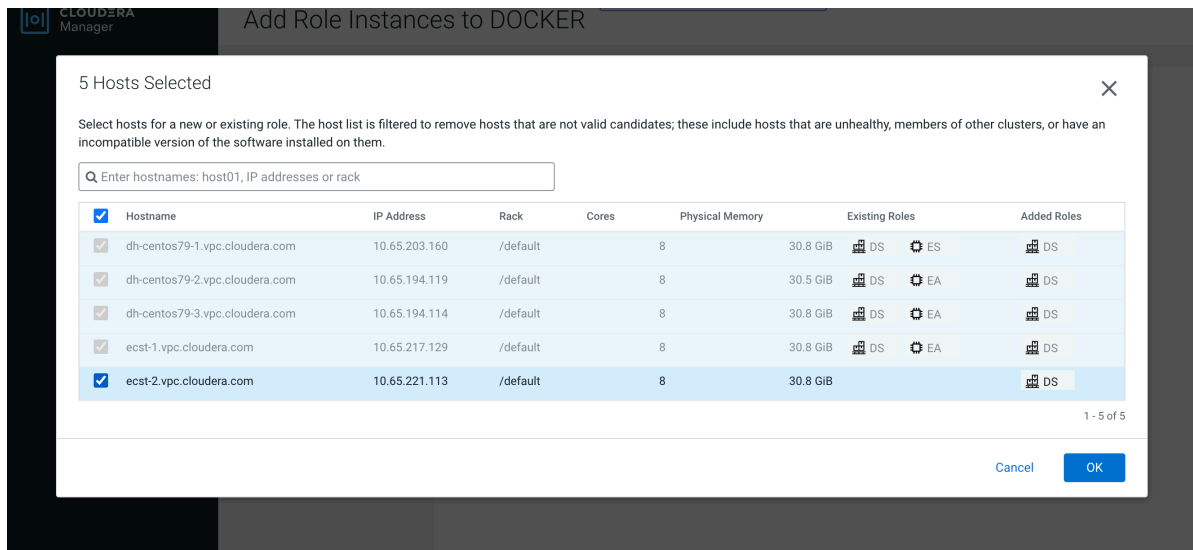
The screenshot shows the Cloudera Manager interface for a cluster named '152-b883'. The 'DOCKER' cluster is selected, and the 'Instances' tab is active. A search bar is present with the text 'Enter search terms (hostname, host ID, IP address, cluster name, rack, health)'. On the left, a sidebar contains navigation links: Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, and Data Services (New). The main content area shows a table of instances with columns: Status, Role Type, Tags, State, Hostname, Commission State, and Role Group. The table lists four instances, all with a status of 'Good Health' and 'Started'. The 'Add Role Instances' button is highlighted in the top right corner.

Status	Role Type	Tags	State	Hostname	Commission State	Role Group
Good Health	Docker Server		Started	dh-centos79-3.vpc.cloudera.com	Commissioned	Docker Server Default Group
Good Health	Docker Server		Started	dh-centos79-1.vpc.cloudera.com	Commissioned	Docker Server Default Group
Good Health	Docker Server		Started	dh-centos79-2.vpc.cloudera.com	Commissioned	Docker Server Default Group
Good Health	Docker Server		Started	ecst-1.vpc.cloudera.com	Commissioned	Docker Server Default Group

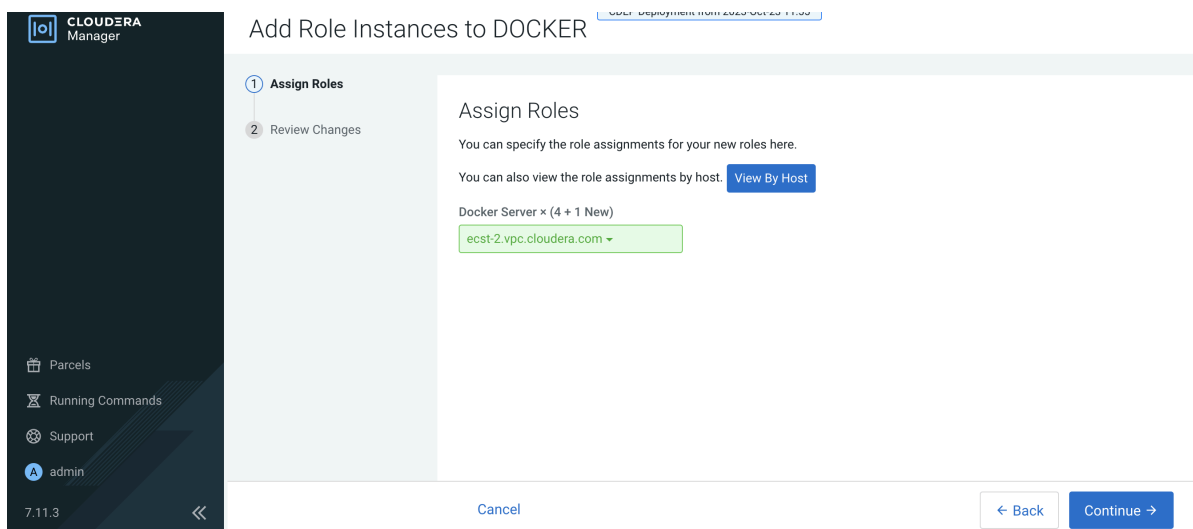
- b. On the Add Role Instances to DOCKER page, click Select hosts.

The screenshot shows the 'Add Role Instances to DOCKER' page. The 'Assign Roles' step is selected, and the 'Select hosts' button is highlighted. The page contains the following text: 'Assign Roles', 'You can specify the role assignments for your new roles here.', 'You can also view the role assignments by host: View By Host', and 'Docker Server x 4'. The 'Select hosts' button is located below the 'Docker Server x 4' text.


- c. On the Hosts Selected pop-up, select the new host, then click OK.



- d. On the Assign Roles page, click Continue.



- e. On the Review Changes page, click Finish.



- f. The new host is listed on the Docker Instances page.

152-b883

CDEP Deployment from 2023-Oct-23 11:25

DOCKER Actions

Status Instances Configuration Commands Charts Library Audits Quick Links

Q Enter search terms (hostname, host ID, IP address, cluster name, rack, health s) Filters Last Updated: Dec 13, 7:00:56 PM UTC

Filters

- STATUS
 - Good Health 4
 - Stopped 1
- COMMISSION STATE
- MAINTENANCE MODE
- RACK ID
- ROLE GROUP
- ROLE TYPE
- STATE
- HEALTH TEST

Actions for Selected

Add Role Instances Role Groups

Status	Role Type	Tags	State	Hostname	Commission State	Role Group
<input type="checkbox"/>	Docker Server		Started	dh-centos79-3.vpc.cloudera.com	Commissioned	Docker Server Default Group
<input type="checkbox"/>	Docker Server		Started	dh-centos79-1.vpc.cloudera.com	Commissioned	Docker Server Default Group
<input type="checkbox"/>	Docker Server		Started	dh-centos79-2.vpc.cloudera.com	Commissioned	Docker Server Default Group
<input type="checkbox"/>	Docker Server		Stopped	ecst-2.vpc.cloudera.com	Commissioned	Docker Server Default Group
<input type="checkbox"/>	Docker Server		Started	ecst-1.vpc.cloudera.com	Commissioned	Docker Server Default Group

1 - 5 of 5

To assign the ECS Agent role group:

- a. Click ECS on the ECS cluster home page, select Instances, then click Add Role Instances.

152-b883

CDEP Deployment from 2023-Oct-23 11:55

ECS Actions

Status Instances Configuration Commands Charts Library Audits Web UI Quick Links

This entity is currently running with an outdated configuration. Restart the service (or the instance) for the changes to take effect.

Q Enter search terms (hostname, host ID, IP address, cluster name, rack, health st) Filters Last Updated: Dec 13, 7:07:48 PM UTC

Filters

- STATUS
 - Good Health 4
- COMMISSION STATE
- MAINTENANCE MODE
- RACK ID
- ROLE GROUP
- ROLE TYPE
- STATE
- HEALTH TEST

Actions for Selected

Add Role Instances Role Groups


Status	Role Type	Tags	State	Hostname	Commission State	Role Group
<input type="checkbox"/>	Ecs Agent		Started	dh-centos79-3.vpc.cloudera.com	Commissioned	Ecs Agent Default Group
<input type="checkbox"/>	Ecs Agent		Started	dh-centos79-2.vpc.cloudera.com	Commissioned	Ecs Agent Default Group
<input type="checkbox"/>	Ecs Agent		Started	ecst-1.vpc.cloudera.com	Commissioned	Ecs Agent Default Group
<input type="checkbox"/>	Ecs Server		Started with Outdated Configuration	dh-centos79-1.vpc.cloudera.com	Commissioned	Ecs Server Default Group

1 - 4 of 4

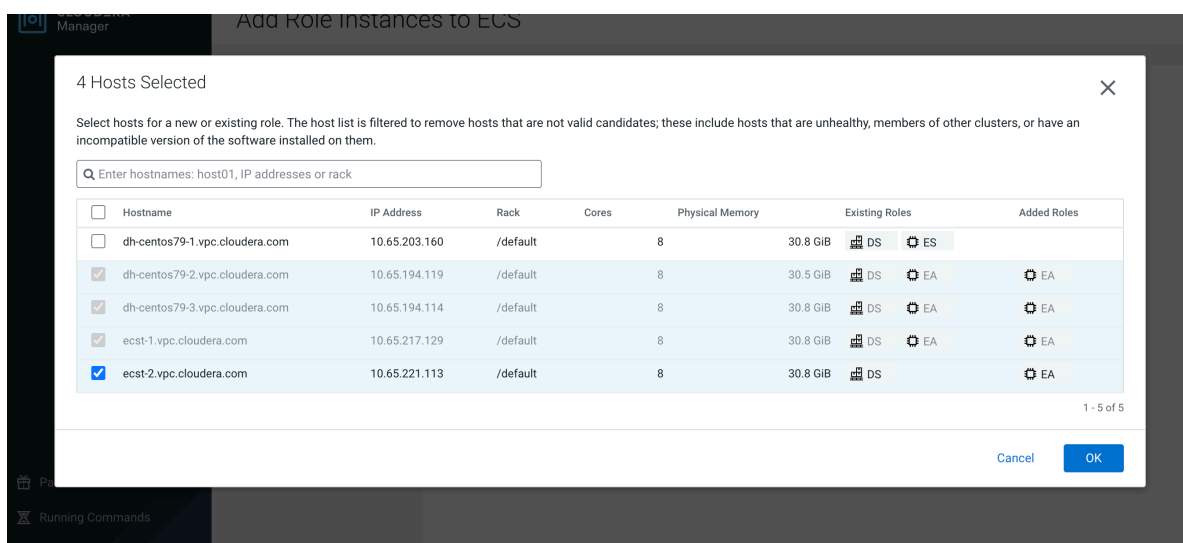
- b. On the Add Role Instances to ECS page, in the Ecs Agent box, click Select hosts.




Important: Be sure to click Select hosts in the Ecs Agent box – do not click the link in the Ecs Server box.




- c. On the Hosts Selected pop-up, select the new host, then click OK.



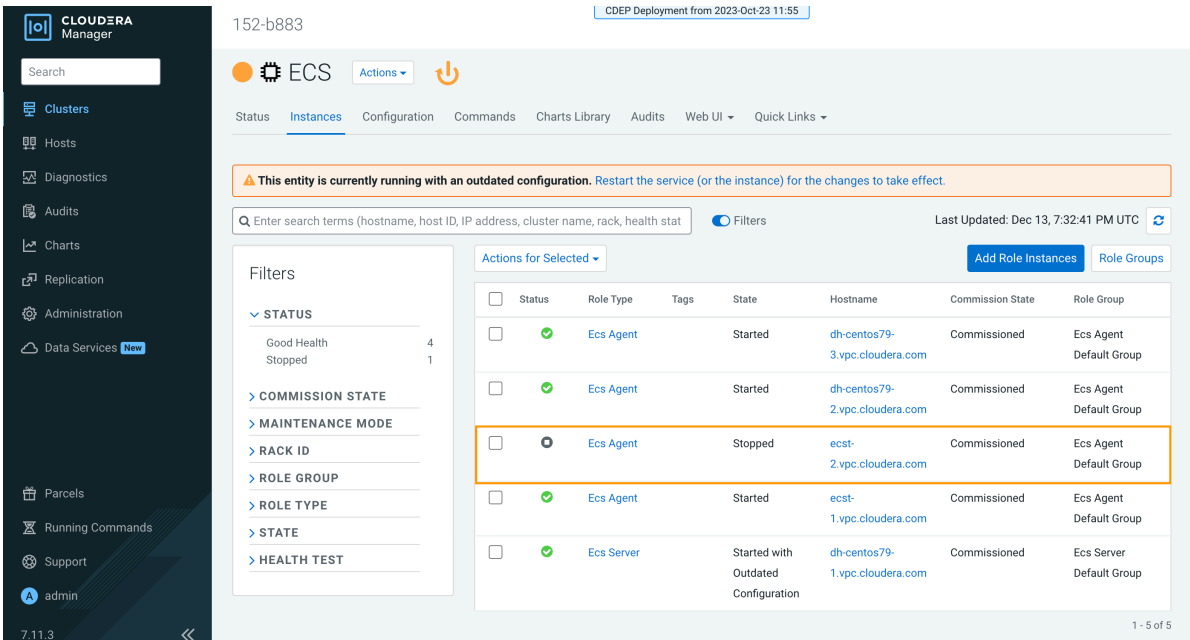
- d. On the Assign Roles page, click Continue.




- e. On the Review Changes page, click Finish.



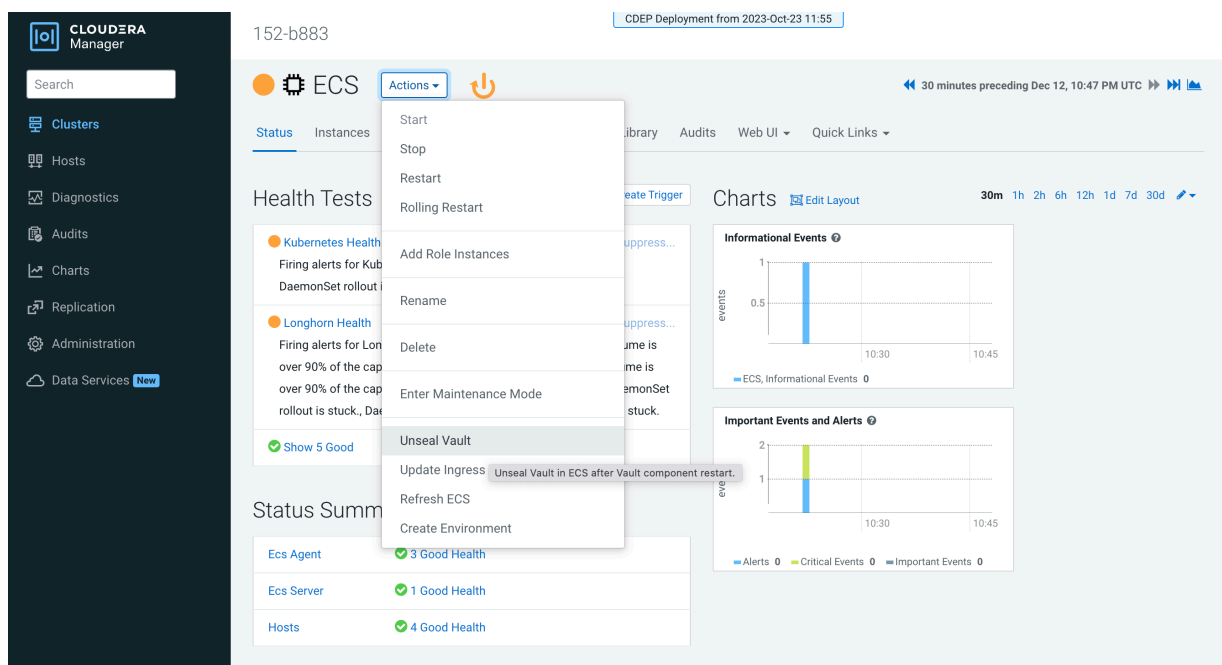
f. The new host is listed on the ECS Instances page.



17. Restart the ECS cluster by clicking the ECS Restart icon, or by selecting Actions > Restart on the ECS cluster home page.



18. Click ECS on the ECS cluster home page, then select Actions > Unseal Vault.



Starting, stopping, restarting, and refreshing Embedded Container Service Clusters

Procedures to start, stop, restart, and refresh Private Cloud Experience clusters

Starting an Embedded Container Service Cluster

Procedure

1. On the Home Status tab, click the Actions Menu to the right of the Embedded Container Service cluster name and select Start.
2. Click the Start button that appears in the next screen to confirm. The Command Details window shows the progress of starting services.

Results

When the All services successfully started message appears, the task is complete and you can close the Command Details window.

Stopping a CDP Private Cloud Data Services Cluster

Procedure

1. On the Home Status tab, click the Actions Menu to the right of the Embedded Container Service cluster name and select Stop.
2. Click the Stop button in the confirmation screen. The Command Details window shows the progress of stopping services.

Results

When the All services successfully stopped message appears, the task is complete and you can close the Command Details window.



Note: The cluster-level Stop action does not stop the Cloudera Management Service. You must stop the Cloudera Management Service separately.

Restarting an Embedded Container Service Cluster

Procedure

1. On the Home Status tab, click the Actions Menu to the right of the cluster name and select Restart.
2. Click the Restart button that appears in the next screen to confirm.
The Command Details window shows the Rolling Restart of services in the cluster. When all the services are restarted successfully, the task is complete and you can close the Command Details window.
3. Click Actions Unseal Vault

Rolling Restart of an Embedded Container Service Cluster

Procedure

1. On the Home Status tab, click the Actions Menu to the right of the cluster name and select Rolling Restart.
2. Click the Rolling Restart button that appears in the next screen to confirm. On this screen, you can select the services (Docker or /and ECS), Roles (Workers only, Non-workers only, All Roles).



Note: Workers only refers to ECS agents, Non-workers only refers to all docker roles and ECS server.

The Command Details window shows the progress of rolling restart of a batch of nodes. Here, batch size refers to the number of worker roles that can be restarted in parallel. The Batch size is 1 by default.

3. Click Actions Unseal Vault

Configuring Restart for an Embedded Container Service cluster

Procedure

1. Navigate to the ECS service, Home Configuration tab.
2. Select Node Readiness Timeout OR Drain Node Timeout configurations to configure overall restart time for an ECS cluster.



Note: Node Readiness Timeout is the maximum time for rescheduling workloads on a new node. This is 5 minutes by default.



Note: Drain Node Timeout is the time out to drain a node. This is 5 minutes by default.

Table 1: ECS Cluster Actions and Performance Impact

ECS Cluster action	Affects Availability	Impact
Stop and Start	Yes	All nodes
Restart	Least	one node at a time
Rolling restart (if the agent batch size is 1 then it is similar to Restart).	Inversely proportional to the batch size.	Depends on batch size.

Refreshing a Embedded Container Service Cluster

Procedure

To refresh a cluster, in the Home Status tab, click the Actions Menu to the right of the cluster name and select Refresh Cluster.



Note:

Refreshing an ECS cluster, runs a cluster refresh action to bring the configuration up to date without restarting all services.

Monitoring Embedded Container Service Clusters

Procedures to monitor Embedded Container Service clusters

Related Information

[Monitoring Services](#)

[Monitoring Clusters](#)

[Docker Server Health Tests](#)

[ECS Health Tests](#)

[ECS Agent Health Tests](#)

[ECS Server Health Tests](#)

[Docker Server Metrics](#)

[ECS Agent Metrics](#)

[ECS Server Metrics](#)

Viewing Health Status

Procedure

1. Open the Cloudera Manager Admin Console.

2. From the Home page, Click on the Embedded Container Service cluster.
3. Click on the ECS or Docker service.

Results

The Service status page displays the Health Test, Status Summary and Health History of the services.

Viewing the Kubernetes Dashboard

About this task

The Kubernetes Dashboard displays configuration and other information about the embedded Kubernetes infrastructure used in the Embedded Container Service cluster. Although you can make configuration changes using the dashboard (if you have the appropriate permissions), you should not make any changes using the dashboard. Cloudera Support may use the dashboard to diagnose problems with the cluster.

Procedure

1. In the Cloudera Manager Admin Console, go to the ECS service.
2. Click Web UI ECS Web UI

Results

The Kubernetes Dashboard displays.

Viewing the Private Cloud Management Console

Procedure

1. In the Cloudera Manager Admin Console, go to the ECS service.
2. Click Web UI Console

Results

The CDP Management Console displays.

Performing maintenance of a single host in the Embedded Container Service cluster

You can perform maintenance on the nodes in your ECS cluster by shutting down the nodes one at a time.

Before you begin

- The containerized cluster must be configured for ECS Server high availability to reduce the downtime.
- You must be able to log into the nodes as root or have sudo privileges.
- The node to be maintained must have a status of Ready. A status of NotReady may suggest the node is having other complicating issues. Run the following command on an ECS server node to verify status of the nodes.

```
/var/lib/rancher/rke2/bin/kubectl --kubeconfig=/etc/rancher/rke2/rke2.yaml  
get nodes
```

Procedure

1. Log in to the Cloudera Manager Admin Console.

2. Stop the ECS role and the Docker server role on the host.
 - Click the Hosts tab.
 - Select the Host->Action->Stop roles on the host.
3. Perform the maintenance on the host.
4. Reboot the host.
5. Log in to the Cloudera Manager Admin Console.
6. Click the Action menu next to the ECS cluster and select Start roles on the host.
7. Click Actions Refresh ECS Cluster.
8. Go to the ECS service page and verify that the Vault is not sealed. This information displays in the Health Tests section.
9. If the Vault is sealed, click Actions Unseal Vault .

Performing the maintenance of all hosts in the Embedded Container Service cluster

If you want to perform the maintenance of all hosts in the ECS cluster follow below steps:

Procedure

1. Log in to the Cloudera Manager Admin Console.
2. Click the Action menu next to the ECS cluster and select Stop.
3. Perform the maintenance on all the hosts.
4. Reboot the hosts.
5. Log in to the Cloudera Manager Admin Console.
6. Click the Action menu next to the ECS cluster and select Start.
7. Click Actions Refresh ECS Cluster .
8. Go to the ECS service page and verify that the Vault is not sealed. This information displays in the Health Tests section.
9. If the Vault is sealed, click Actions Unseal Vault .

Configuring a containerized cluster with SELinux

You can configure a containerized cluster with SELinux to enable it to run the Embedded Container Service (ECS).

Procedure

1. Ensure that the hosts you use for the containerized cluster meet all [hardware](#) and [software](#) requirements for use with CDP Private Cloud Data Services.
2. Enable SELinux in Permissive mode by updating the `/etc/selinux/config` file on all ECS hosts by running the following commands:

```
sed -i 's/SELINUX=disabled/SELINUX=permissive/' /etc/selinux/config
reboot
```

3. Add the SELinux policies provided by RKE2 by installing the RPMs on all ECS hosts. Use the following commands:

```
yum localinstall -y http://mirror.centos.org/centos/7/extras/x86_64/Packages/container-selinux-2.107-3.el7.noarch.rpm
wget https://github.com/rancher/rke2-selinux/releases/download/v0.8.stable.2/rke2-selinux-0.8-2.el7.noarch.rpm
yum install -y rke2-selinux-0.8-2.el7.noarch.rpm
```

4. Uninstall the nscd service by running the following command on all ECS hosts :

```
yum erase -y nscd
```

5. Install a containerized cluster on all hosts. See [Adding a CDP Private Cloud Data Services cluster](#).
6. Enable SELinux in Enforced mode by running the following commands on all ECS hosts:

```
setenforce 1
```

You can confirm that SELinux is running in Enforced mode by running the following command:

```
getenforce
```

7. Verify that the ECS cluster hosts are sending heartbeats to the Cloudera Manager server.
 - a) Open the Cloudera Manager Admin Console.
 - b) Click Hosts All Hosts .
 - c) Check the Last Heartbeat column for heartbeat status.
8. Verify that your workloads are functioning as expected.

Decommissioning ECS Hosts

You can decommission ECS hosts and remove them from the cluster.

About this task

1. Cordon the node. Longhorn will automatically disable the node scheduling when a Kubernetes node is cordoned. Run the following command on any ECS Server host:

```
kubectl cordon [***node***]
```

2. Drain the node to move the workload to somewhere else. Run the following command on any ECS Server host:

```
kubectl drain [***node***] --ignore-daemonsets --pod-selector='app!=csi-attacher,app!=csi-provisioner' --delete-emptydir-data
```

3. Detach all the volumes on the node. Navigate to the ECS Service page on Cloudera Manager UI.
 - a. In the Web UI dropdown, select Storage UI to open the Longhorn UI.
 - b. Under the Volume tab in Longhorn UI, select the volumes on this node. Click Detach and select Yes on the screen prompt.

If the node has been drained, all the workloads should be migrated to another node already.

If there are any other volumes remaining attached, detach them before continuing.

4. Remove the node from Longhorn using the Delete in the Node tab. Or, remove the node from Kubernetes. Run the following command on any ECS Server host:

```
kubectl delete node [***node-name***]
```

Longhorn will automatically remove the node from the cluster.

5. Uninstall ECS and Docker artifacts from the host. Run below commands on the host:

```
cd /opt/cloudera/parcels/ECS/bin
./rke2-killall.sh # usually 2 times is sufficient
./rke2-uninstall.sh
rm -rf /ecs/* # assumes the default defaultDataPath and IsoDataPath
rm -rf /var/lib/docker_server/* # deletes the auth and certs
rm -rf /etc/docker/certs.d/* # delete the ca.crt
```

```
rm -rf /docker # assumes the default defaultDataPath for docker
```

- Go to the Hosts page for the ECS Cluster, select that host, and under Actions for Selected, click Begin Maintenance (Suppress Alerts/Decommission)

Dedicating ECS nodes for specific workloads

You use Cloudera Manager to dedicate Embedded Container Service (ECS) cluster nodes for specific workloads. You can dedicate GPU nodes for CML workloads, and NVME nodes for CDW workloads.

Dedicating ECS nodes when creating a new cluster

- Check the [ECS installation requirements](#).
- [Add the new hosts to Cloudera Manager](#).
- In Cloudera Manager, click Hosts > All Hosts, then select one or more of the new ECS hosts.
- Click the Configuration tab, then use the Search box to locate the `node_taint` configuration property.
- Select Dedicated GPU Node to dedicate the node for CML workloads, or select Dedicated NVME node to dedicate the node for CDW workloads.

When either of these options are selected, no other workload pods will be allowed to run on the dedicated node.

The screenshot shows the Cloudera Manager interface. On the left is a sidebar with navigation links: Clusters, Hosts (selected), Diagnostics, Audits, Charts, Replication, Administration, Data Services (New), Parcels, Running Commands, Support, and a user profile 'admin'. The main panel is titled 'Hosts Configuration' and shows a search for 'node_taint'. On the left, a 'Filters' sidebar shows counts for SCOPE (All Hosts: 1), CATEGORY (Advanced: 1, Monitoring: 0, Parcels: 0, Resource Management: 0), and STATUS (Error: 0, Warning: 0, Edited: 1, Non-Default: 1, Include Overrides: 0). The main configuration area shows 'Data Services: Restrict workloads types' with a radio button selected for 'Dedicated GPU Node'. Other options are 'Dedicated NVME Node' and 'None'. There are links for 'Undo', 'Add Host Overrides', and 'Show All Descriptions'. At the bottom, it says '1 Edited Value' and 'Reason for change: Modified Data Services: Restrict workloads types', with a 'Save Changes(CTRL+S)' button.

- Click Save Changes.
- Repeat the previous steps to add the other ECS hosts to Cloudera Manager and assign workload types.
- Follow the [ECS installation procedure](#). When you reach the Specify Hosts page in the installation wizard, the hosts you added to Cloudera Manager appear. Select the hosts, click Continue, then proceed through the rest of the installation wizard.
- After the installation is complete, the applicable workloads will only run on the specified dedicated nodes.

Dedicating ECS nodes in an existing cluster

- Open the Cloudera Manager Admin Console.
- On the Home page, click the ECS Cluster.

3. Click Hosts, select one or more of the ECS hosts, then click the Configuration tab.
4. Click the Configuration tab, then use the Search box to locate the `node_taint` configuration property.
5. Select Dedicated GPU Node to dedicate the node for CML workloads, or select Dedicated NVME node to dedicate the node for CDW workloads.

When either of these options are selected, no other workload pods will be allowed to run on the dedicated node.

The screenshot shows the Cloudera Manager interface. On the left is a dark sidebar with navigation links: Clusters, Hosts (selected), Diagnostics, Audits, Charts, Replication, Administration, Data Services (marked as New), Parcels, Running Commands, Support, and a user profile for 'admin'. The main area is titled 'Hosts Configuration' and has a search bar containing 'node_taint'. Below the search bar are three filter sections: SCOPE (All Hosts: 1), CATEGORY (Advanced: 1, Monitoring: 0, Parcels: 0, Resource Management: 0), and STATUS (Error: 0, Warning: 0, Edited: 1, Non-Default: 1, Include Overrides: 0). The configuration details for 'node_taint' are shown on the right, with the option 'Dedicated GPU Node' selected under 'Data Services: Restrict workloads types'. Other options include 'Dedicated NVME Node' and 'None'. There are links for 'Undo', 'Add Host Overrides', and 'Show All Descriptions'. At the bottom, a status bar indicates '1 Edited Value' and 'Reason for change: Modified Data Services: Restrict workloads types', with a 'Save Changes(CTRL+S)' button.

6. Click Save Changes.
7. Repeat the previous steps to assign workload types to the other ECS hosts.
8. On the ECS Cluster landing page, click Actions > Refresh Cluster.
9. After the Refresh is complete, click Actions > Rolling Restart.

Specifying racks for ECS clusters

You use Cloudera Manager to assign Embedded Container Service (ECS) cluster hosts to a specific rack.

About this task

- All hosts in an ECS cluster must have the same assigned rack name and path structure. A configuration error will occur if the rack names do not match.
- ECS cluster hosts with no specified rack name are assigned the default rack name value. The default value means that no rack name has been specified for the ECS cluster hosts.

Specifying a rack name for an ECS cluster

1. In Cloudera Manager, select the ECS cluster, then click Hosts.

2. In the Hosts list, click the top checkbox to select all of the cluster hosts.

CLOUDERA
Manager

Search

Clusters

Hosts

Diagnostics

Audits

Charts

Replication

Administration

Data Services New

152-b813

CDEP Deployment from 2023-Sep-26 08:29

Cust

Hosts

Configuration

Add Hosts

Review Upgrade Status

Inspect Hosts in Cluster

Inspect Cluster Network Performance

Search

Filters

Last Updated: Oct 1, 7:41:54 PM UTC

Columns: 11 Selected

Filters

STATUS

Good Health 3

CLUSTERS

CORES

COMMISSION STATE

LAST HEARTBEAT

LOAD (1 MINUTE)

LOAD (5 MINUTES)

LOAD (15 MINUTES)

MAINTENANCE MODE

UPGRADE DOMAIN

RACK

SERVICE

Actions for Selected (3)

Status	Name	IP	Roles	Tags	Commission Stat
<input checked="" type="checkbox"/>	dh-centos79-1.vpc.cloudera.com	10.65.201.209	2 Roles		Commissioned
<input checked="" type="checkbox"/>	dh-centos79-2.vpc.cloudera.com	10.65.194.34	2 Roles		Commissioned
<input checked="" type="checkbox"/>	dh-centos79-3.vpc.cloudera.com	10.65.200.38	2 Roles		Commissioned

1 - 3 of 3

3. Click Actions for Selected, then click Assign Rack.

CLOUDERA
Manager

Search

Clusters

Hosts

Diagnostics

Audits

Charts

Replication

Administration

Data Services New

152-b813

CDEP Deployment from 2023-Sep-26 08:29

Cust

Hosts

Configuration

Add Hosts

Review Upgrade Status

Inspect Hosts in Cluster

Inspect Cluster Network Performance

Search

Filters

Last Updated: Oct 1, 7:47:54 PM UTC

Columns: 11 Selected

Filters

STATUS

Good Health 3

CLUSTERS

CORES

COMMISSION STATE

LAST HEARTBEAT

LOAD (1 MINUTE)

LOAD (5 MINUTES)

LOAD (15 MINUTES)

MAINTENANCE MODE

UPGRADE DOMAIN

RACK

SERVICE

Actions for Selected (3)

Assign Rack

Assign Upgrade Domain

Regenerate Keytab

Apply Host Template

Start Roles on Hosts

Stop Roles on Hosts

Begin Maintenance (Suppress Alerts/Decommission)

End Maintenance (Enable Alerts/Recommission)

Edit Tags

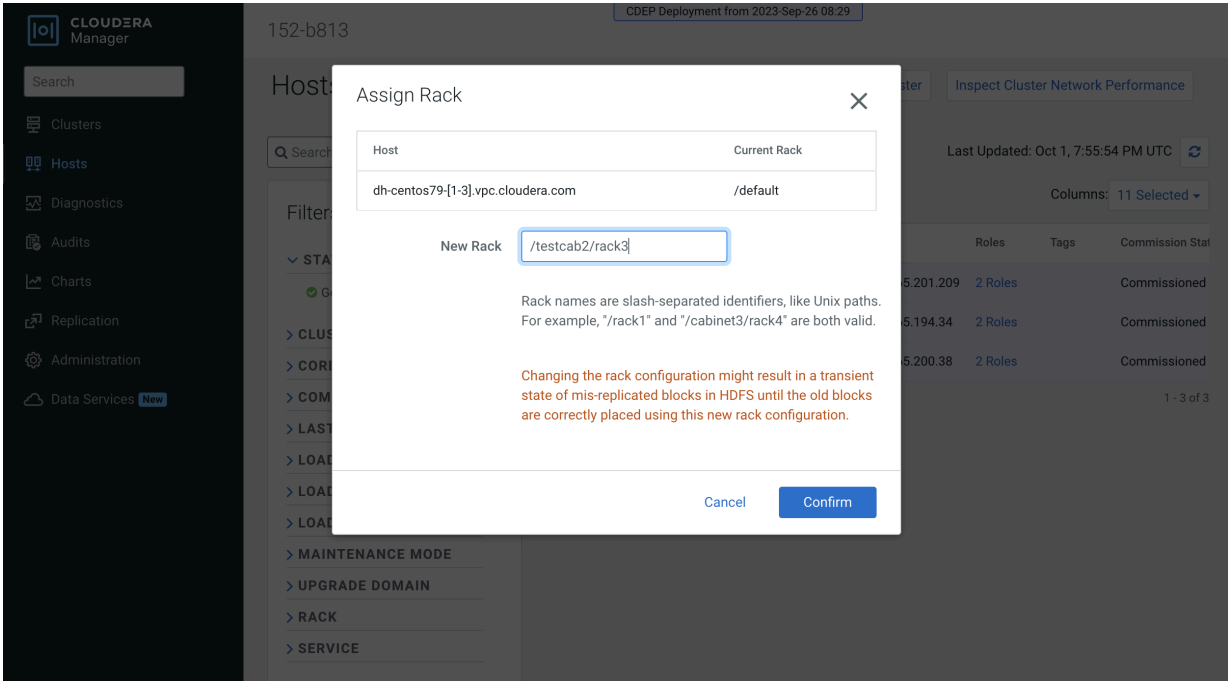
Remove From Cluster

Remove From Cloudera Manager

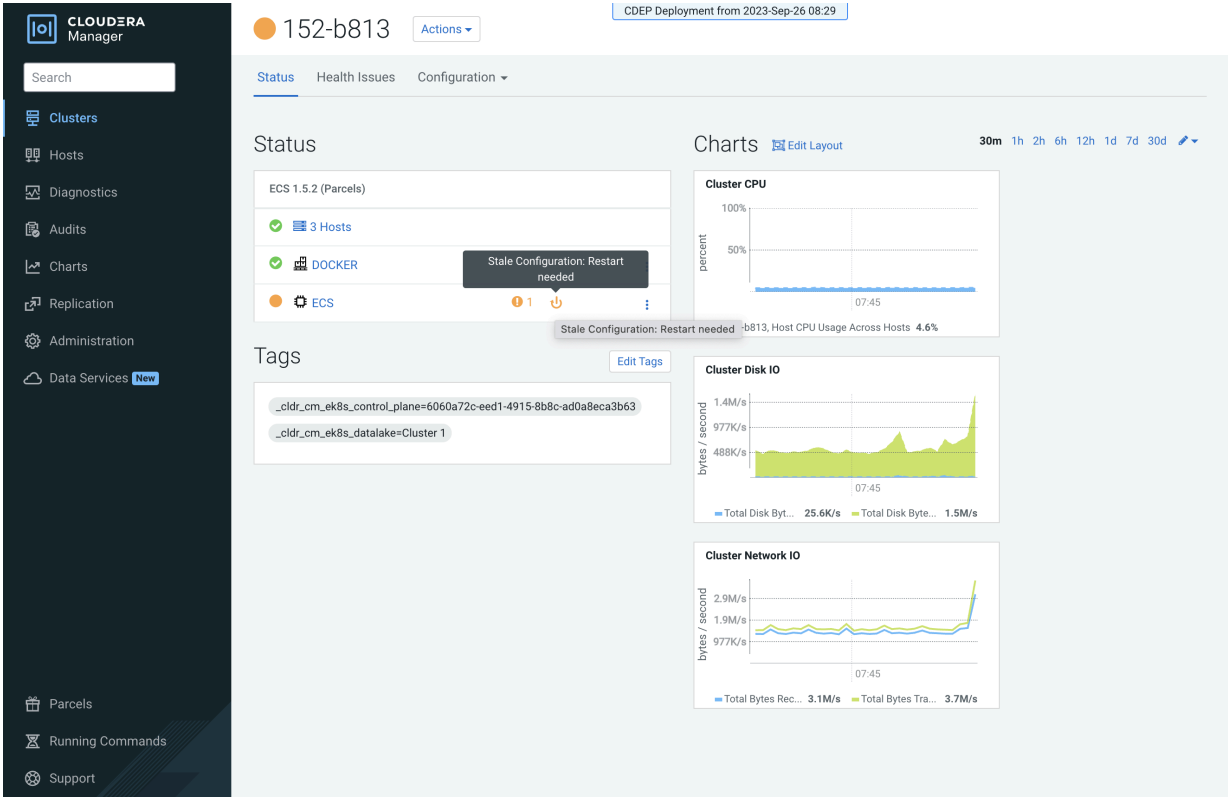
Roles	Tags	Commission Stat
201.209 2 Roles		Commissioned
.194.34 2 Roles		Commissioned
.200.38 2 Roles		Commissioned

1 - 3 of 3

4. On the Assign Rack popup, enter a rack name in the New Rack box, then click Confirm.



5. Cloudera Manager detects this configuration change, and displays a Stale Configuration warning. You must restart the cluster in order for the updated configuration to take effect.



6. Click the Stale Configuration icon, then click Restart Stale Services and click through the Restart wizard.

7. When the Restart is complete, you can use the Assign Rack popup to confirm that the new rack name has been applied to the ECS cluster hosts.

- You can also use the ECS Web UI to view cluster host rack assignments. Select the ECS cluster, click ECS, then click Web UI > ECS Web UI . In the Web UI, select the CDP namespace, then click Nodes.

Note that in Kubernetes periods are used as separators (rather than slashes) in the rack name path. The leading slash is also not used in Kubernetes.

Name	Labels	CPU Ready requests (cores)	CPU limits (cores)	CPU capacity (cores)	Memory requests (bytes)	Memory limits (bytes)	Memory capacity (bytes)	Pods
dh-centos79-3.vpc.cloudera.com	beta.kubernetes.io/arch: amd64 beta.kubernetes.io/os: linux kubernetes.io/arch: amd64 kubernetes.io/hostname: dh-centos79-3.vpc.cloudera.com kubernetes.io/os: linux rack: testcab2.rack3	True	6.07 (75.81%)	6.95 (86.88%)	8.00	8.82Gi (28.61%)	29.13Gi (94.54%)	30.81Gi 39 (7.8%)
dh-centos79-2.vpc.cloudera.com	beta.kubernetes.io/arch: amd64 beta.kubernetes.io/os: linux kubernetes.io/arch: amd64 kubernetes.io/hostname: dh-centos79-2.vpc.cloudera.com kubernetes.io/os: linux rack: testcab2.rack3	True	7.92 (99.01%)	7.55 (94.38%)	8.00	13.78Gi (45.21%)	28.98Gi (95.07%)	30.48Gi 48 (9.6%)
dh-centos79-1.vpc.cloudera.com	beta.kubernetes.io/arch: amd64 beta.kubernetes.io/os: linux ecs_role: master kubernetes.io/arch: amd64 kubernetes.io/hostname: dh-centos79-1.vpc.cloudera.com kubernetes.io/os: linux node-role.kubernetes.io/control-plane: true node-role.kubernetes.io/etcd: true node-role.kubernetes.io/master: true rack: testcab2.rack3	True	7.97 (99.63%)	11.35 (141.88%)	8.00	11.36Gi (36.88%)	29.85Gi (96.90%)	30.81Gi 57 (11.40%)

Specifying a rack name when creating a new ECS cluster

Currently the ECS installation wizard does not enable you to assign rack names when creating a new ECS cluster. Therefore, you should first add the new set of ECS hosts to Cloudera Manager, and then assign the rack name in Cloudera Manager. You can then use the ECS installation wizard to create a new ECS cluster using these hosts.

- Check the [ECS installation requirements](#).
- [Add the new hosts to Cloudera Manager](#).

3. In Cloudera Manager, click Hosts > All Hosts, then select the hosts you just added.

CLUSTERA

Manager

Search

Clusters

Hosts

Diagnostics

Audits

Charts

Replication

Administration

Data Services New

Home

CDEP Deployment from 2023-Sep-26 08:29

All Hosts

ConfigurationAdd HostsReview Upgrade StatusInspect All HostsInspect Network Performance

Search

Filters

Last Updated: Oct 2, 8:03:05 PM UTC

Filters

Actions for Selected (3)

Columns: 11 Selected

	Status	Name	IP	Roles	Tags	Commission Statu
<input type="checkbox"/>	✓	dh-centos79-1.vpc.cloudera.com	10.65.201.209	2 Roles		Commissioned
<input type="checkbox"/>	✓	dh-centos79-2.vpc.cloudera.com	10.65.194.34	2 Roles		Commissioned
<input type="checkbox"/>	✓	dh-centos79-3.vpc.cloudera.com	10.65.200.38	2 Roles		Commissioned
<input checked="" type="checkbox"/>	✓	dh-centos79t-1.vpc.cloudera.com	10.65.199.15			Commissioned
<input checked="" type="checkbox"/>	✓	dh-centos79t-2.vpc.cloudera.com	10.65.205.101			Commissioned
<input checked="" type="checkbox"/>	✓	dh-centos79t-3.vpc.cloudera.com	10.65.200.0			Commissioned
<input type="checkbox"/>	✓	dhoyle711318-1.dhoyle711318.root.hwx.site	172.27.173.77	55 Roles		Commissioned
<input type="checkbox"/>	✓	dhoyle711318-2.dhoyle711318.root.hwx.site	172.27.76.66	23 Roles		Commissioned
<input type="checkbox"/>	✓	dhoyle711318-3.dhoyle711318.root.hwx.site	172.27.203.76	18 Roles		Commissioned

1 - 9 of 9

4. Click Actions for Selected, then click Assign Rack.

CLUSTERA

Manager

Search

Clusters

Hosts

Diagnostics

Audits

Charts

Replication

Administration

Data Services New

152-b813

CDEP Deployment from 2023-Sep-26 08:29

Hosts

ConfigurationAdd HostsReview Upgrade StatusInspect Hosts in ClusterInspect Cluster Network Performance

Search

Filters

Last Updated: Oct 1, 7:47:54 PM UTC

Filters

Actions for Selected (3)

Columns: 11 Selected

	Roles	Tags	Commission Statu
201.209	2 Roles		Commissioned
194.34	2 Roles		Commissioned
200.38	2 Roles		Commissioned

1 - 3 of 3

Assign Rack

Assign Upgrade Domain

Regenerate Keytab

Apply Host Template

Start Roles on Hosts

Stop Roles on Hosts

Begin Maintenance (Suppress Alerts/Decommission)

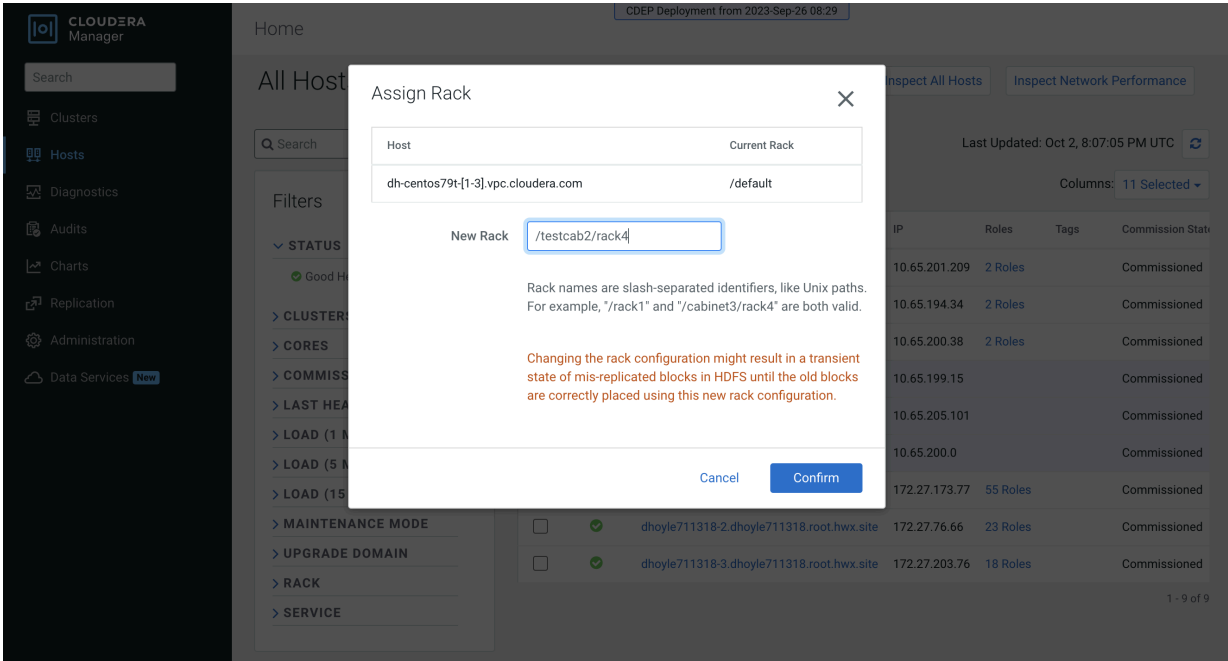
End Maintenance (Enable Alerts/Recommission)

Edit Tags

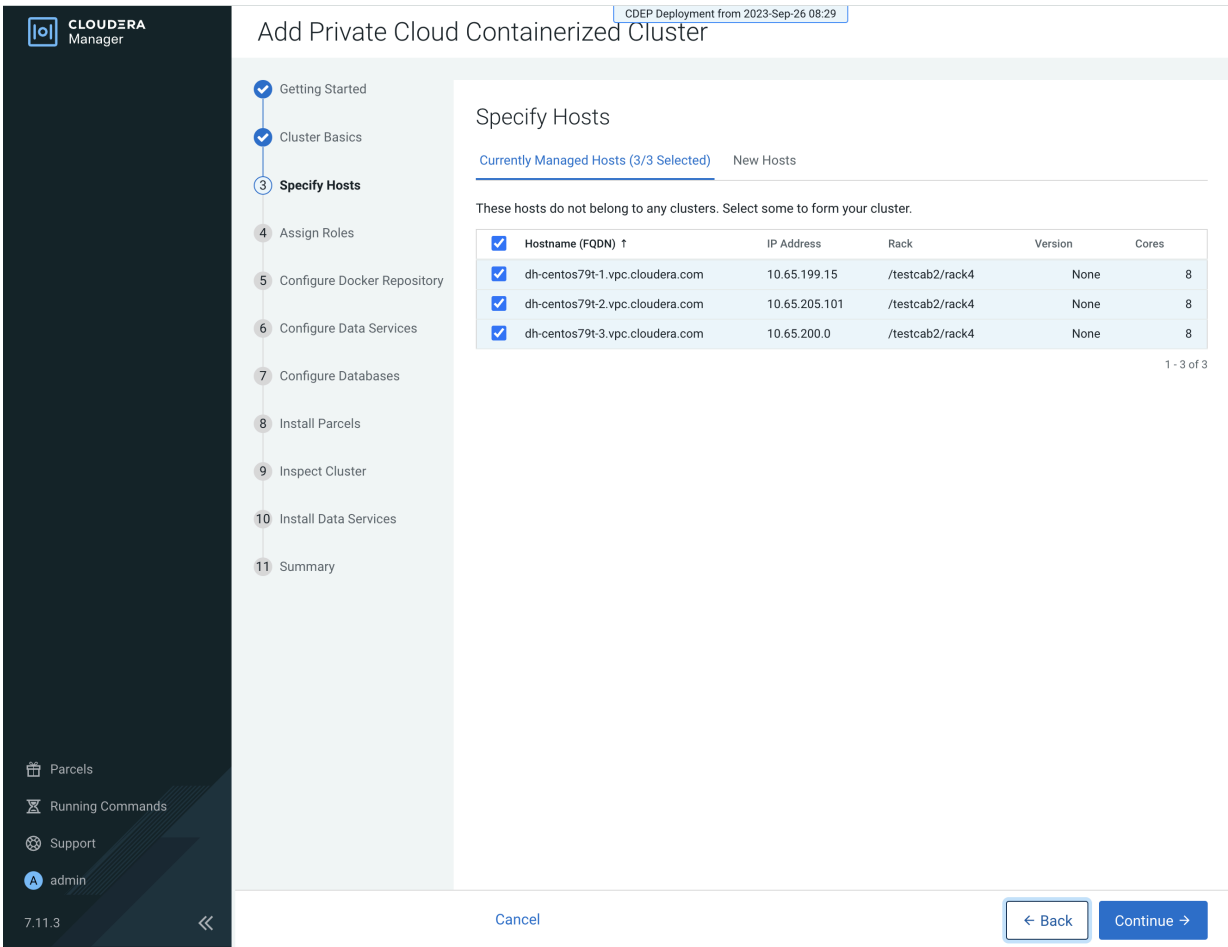
Remove From Cluster

Remove From Cloudera Manager

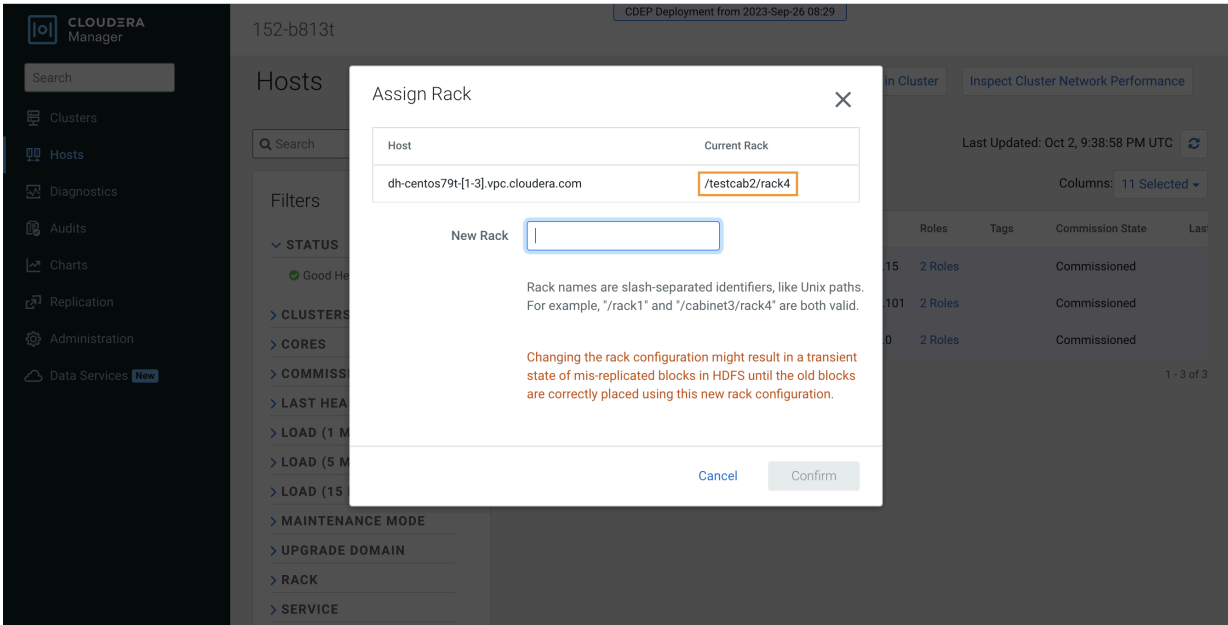
5. On the Assign Rack popup, enter a rack name in the New Rack box, then click Confirm.



6. Follow the [ECS installation procedure](#). When you reach the Specify Hosts page in the installation wizard, the hosts you added to Cloudera Manager appear. Select the hosts, click Continue, then proceed through the rest of the installation wizard.



- 7. After the installation is complete, you can use the Assign Rack popup or the ECS Web UI to view the rack assignments for the ECS cluster hosts.



<div> kubernetes </div> <div> <div>cdp</div> <div>Search</div> </div> <div> </div>										
Cluster > Nodes										
<div> <div>Daemon Sets</div> <div>Deployments</div> <div>Jobs</div> <div>Pods</div> <div>Replica Sets</div> <div>Replication Controllers</div> <div>Stateful Sets</div> <div>Service</div> <div>Ingresses N</div> <div>Ingress Classes</div> <div>Services N</div> <div>Config and Storage</div> <div>Config Maps N</div> <div>Persistent Volume Claims N</div> <div>Secrets N</div> <div>Storage Classes</div> <div>Cluster</div> <div>Cluster Role Bindings</div> <div>Cluster Roles</div> <div>Events N</div> <div>Namespaces</div> <div>Network Policies N</div> <div>Nodes</div> <div>Persistent Volumes</div> <div>Role Bindings N</div> <div>Roles N</div> <div>Service Accounts N</div> <div>Custom Resource Definitions</div> <div>Settings</div> <div>About</div> </div>										
Nodes										
	Name	Labels	Ready	CPU requests (cores)	CPU limits (cores)	CPU capacity (cores)	Memory requests (bytes)	Memory limits (bytes)	Memory capacity (bytes)	Pods
	● dh-centos79t-2.vpc.cloudera.com	beta.kubernetes.io/arch: amd64	True	7.44 (93.03%)	5.10 (63.75%)	8.00	26.18Gi (85.88%)	21.64Gi (70.99%)	30.48Gi	49 (9.80%)
		beta.kubernetes.io/os: linux								
		kubernetes.io/arch: amd64								
		kubernetes.io/hostname: dh-centos79t-2.vpc.cloudera.com								
	● dh-centos79t-3.vpc.cloudera.com	kubernetes.io/os: linux	True	7.62 (95.26%)	8.35 (104.38%)	8.00	10.48Gi (34.40%)	36.83Gi (120.83%)	30.48Gi	52 (10.40%)
		rack: testcab2.rack4								
		Show less								
		beta.kubernetes.io/arch: amd64								
	● dh-centos79t-1.vpc.cloudera.com	beta.kubernetes.io/os: linux	True	6.40 (79.94%)	9.40 (117.50%)	8.00	8.91Gi (28.93%)	25.66Gi (83.30%)	30.81Gi	47 (9.40%)
		ecs_role: master								
		kubernetes.io/arch: amd64								
		kubernetes.io/hostname: dh-centos79t-1.vpc.cloudera.com								
		kubernetes.io/os: linux	True	6.40 (79.94%)	9.40 (117.50%)	8.00	8.91Gi (28.93%)	25.66Gi (83.30%)	30.81Gi	47 (9.40%)
		node-role.kubernetes.io/control-plane: true								
		node-role.kubernetes.io/etcd: true								
		node-role.kubernetes.io/master: true								
		node-role.kubernetes.io/master: true	True	6.40 (79.94%)	9.40 (117.50%)	8.00	8.91Gi (28.93%)	25.66Gi (83.30%)	30.81Gi	47 (9.40%)
		node-role.kubernetes.io/master: true								
		node-role.kubernetes.io/master: true								
		node-role.kubernetes.io/master: true								
		node-role.kubernetes.io/master: true	True	6.40 (79.94%)	9.40 (117.50%)	8.00	8.91Gi (28.93%)	25.66Gi (83.30%)	30.81Gi	47 (9.40%)
		node-role.kubernetes.io/master: true								
		node-role.kubernetes.io/master: true								
		node-role.kubernetes.io/master: true								


Adding a host to an ECS cluster with a previously specified rack name

When you add a host directly to an ECS cluster, there is no way to specify a rack name for the new host, so it will be assigned the default rack name. A configuration error will occur if you try to add a new host directly to an ECS cluster with a previously specified rack name, since the default rack name of the new host does not match the rack name previously assigned to the other cluster hosts.


Therefore, you should first add the new ECS host to Cloudera Manager, and then use Cloudera Manager to assign the same rack name as the other ECS cluster hosts to the new host. You can then add the new host to the ECS cluster.

1. Check the [ECS installation requirements](#).

2. [Add the new hosts to Cloudera Manager](#). You can also access the Add Hosts wizard by clicking Hosts in the ECS cluster, and then clicking Add Hosts. Select Add hosts to Cloudera Manager.



3. In Cloudera Manager, click Hosts, then select the host you just added.



4. Click Actions for Selected, then click Assign Rack.


The screenshot shows the Cloudera Manager interface. On the left is a sidebar with navigation links: Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, and Data Services (marked as 'New'). The main area is titled 'All Hosts' and includes a search bar, a filter toggle, and a 'Last Updated' timestamp. A dropdown menu 'Actions for Selected (1)' is open, listing various actions. 'Assign Rack' is highlighted. In the background, a table of hosts is visible with columns: IP, Roles, Tags, and Commission State. The table shows 10 hosts, all with a 'Commissioned' state.

IP	Roles	Tags	Commission State
10.65.201.209	2 Roles		Commissioned
10.65.194.34	2 Roles		Commissioned
10.65.200.38	2 Roles		Commissioned
10.65.192.56			Commissioned
10.65.199.15	2 Roles		Commissioned
10.65.205.101	2 Roles		Commissioned
10.65.200.0	2 Roles		Commissioned
172.27.173.77	55 Roles		Commissioned
172.27.76.66	23 Roles		Commissioned
172.27.203.76	18 Roles		Commissioned

5. On the Assign Rack popup, enter the same rack name assigned to the other ECS cluster hosts in the New Rack box, then click Confirm.

The screenshot shows the 'Assign Rack' dialog box. It has a 'Host' field with the value 'dh-centos79a-1.vpc.cloudera.com' and a 'Current Rack' field with the value '/default'. Below these is a 'New Rack' field containing '/testcab2/rack3'. A warning message states: 'Rack names are slash-separated identifiers, like Unix paths. For example, "/rack1" and "/cabinet3/rack4" are both valid.' Another warning in red text says: 'Changing the rack configuration might result in a transient state of mis-replicated blocks in HDFS until the old blocks are correctly placed using this new rack configuration.' At the bottom are 'Cancel' and 'Confirm' buttons.

6. In the ECS cluster, click Hosts, then click Add Hosts. Select Add hosts to Cluster, then click Continue.

 CLOUDERA
Manager

Parcels

Running Commands

Support

admin

7.11.3

Add Hosts

CDEP Deployment from 2023-Sep-26 08:29

The Add Hosts Wizard allows you to install the Cloudera Manager Agent on new hosts. You can either keep the new hosts available to be added to a cluster in the future, or you can add new hosts to an existing cluster

☐ Add hosts to Cloudera Manager

You can use these hosts later to create new clusters or expand existing clusters.

☒ Add hosts to Cluster

152-b813

← Back

Continue →

7. On the Specify Hosts page, select the new host, then click through the rest of the Add Hosts wizard.

CloudERA
Manager

Parcels

Running Commands

Support

admin

7.11.3

<<

1 Specify Hosts

2 Install Parcels

3 Inspect Hosts

4 Select Host Template

5 Deploy Client Config

Add Hosts

CDEP Deployment from 2023-Sep-26 08:29

Specify Hosts

Currently Managed Hosts (1/1 Selected) New Hosts

These hosts do not belong to any clusters. Select some to form your cluster.

<input checked="" type="checkbox"/>	Hostname (FQDN) ↑	IP Address	Rack	Version	Cores
<input checked="" type="checkbox"/>	dh-centos79a-1.vpc.cloudera.com	10.65.192.56	/testcab2/rack3	None	8

1 - 1 of 1

Cancel

< Back

Continue >

8. After the Add Host wizard is completed, the new host appears on the ECS cluster Hosts page.

CloudERA
Manager

Search

Clusters

Hosts

Diagnostics

Audits

Charts

Replication

Administration

Data Services New

152-b813

CDEP Deployment from 2023-Sep-26 08:29

Hosts

Configuration

Add Hosts

Review Upgrade Status

Inspect Hosts in Cluster

Inspect Cluster Network Performance

Search

Filters

Last Updated: Oct 3, 6:56:46 PM UTC

Filters

STATUS

Good Health 4

CLUSTERS

CORES

COMMISSION STATE

LAST HEARTBEAT

LOAD (1 MINUTE)

LOAD (5 MINUTES)

LOAD (15 MINUTES)

MAINTENANCE MODE

UPGRADE DOMAIN

RACK

SERVICE

Actions for Selected (1)

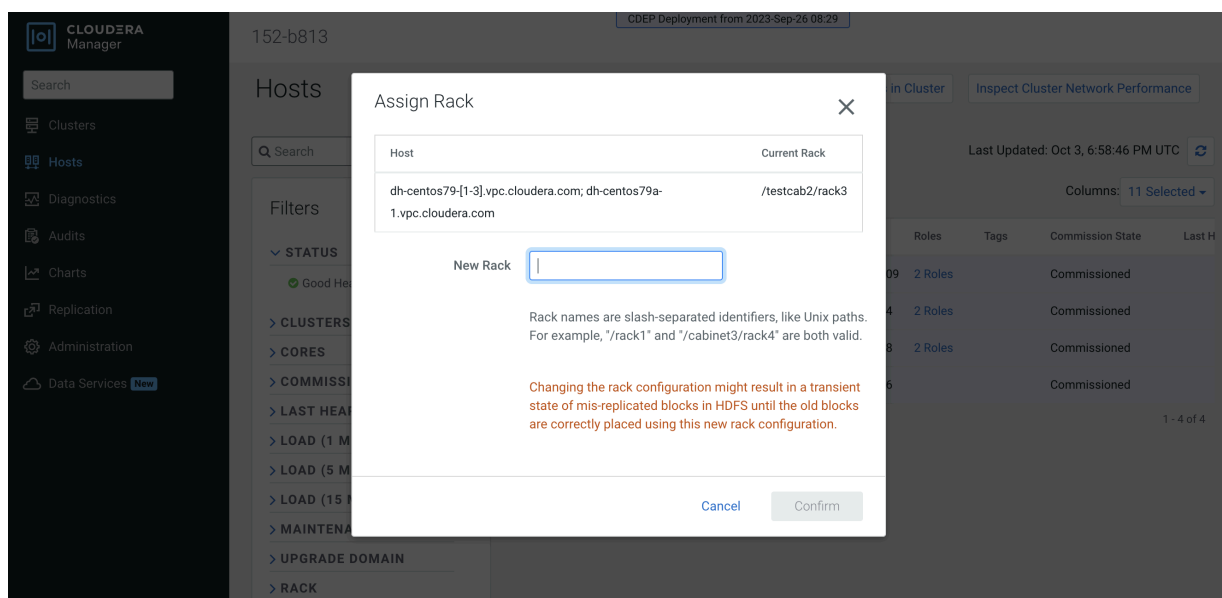
Columns: 11 Selected

<input type="checkbox"/>	Status	Name	IP	Roles	Tags	Commission State	Last H
<input type="checkbox"/>	✓	dh-centos79-1.vpc.cloudera.com	10.65.201.209	2 Roles		Commissioned	
<input type="checkbox"/>	✓	dh-centos79-2.vpc.cloudera.com	10.65.194.34	2 Roles		Commissioned	
<input type="checkbox"/>	✓	dh-centos79-3.vpc.cloudera.com	10.65.200.38	2 Roles		Commissioned	
<input checked="" type="checkbox"/>	✓	dh-centos79a-1.vpc.cloudera.com	10.65.192.56			Commissioned	

1 - 4 of 4

41

9. You can use the Assign Rack popup to view the rack assignments for the ECS cluster hosts and confirm that the rack name for the new host matches the rack name of the other cluster hosts.



ECS unified time zone

You can synchronize the Embedded Container Service (ECS) cluster time zone with the Cloudera Manager Base time zone.

In CDP Private Cloud Data Services versions earlier than 1.5.2, containers running on an ECS Kubernetes cluster did not inherit the time zone settings from the Cloudera Manager Base host. In most cases, Kubernetes containers use Coordinated Universal Time (UTC) by default.

In Private Cloud Data Services 1.5.2 and higher versions, you can unify the time zone in the ECS cluster with the Cloudera Manager Base time zone. All workload pods in the ECS cluster run under the Cloudera Manager time zone, and workload logs on the ECS cluster are correlated with the Cloudera Manager Base logs. Timestamp-related SQL queries are also correlated.

- Unified time zone is enabled by default for new CDP Private Cloud Data Services 1.5.2+ installs.
- When upgrading from earlier versions of CDP Private Cloud Data Services to 1.5.2+, unified time zone is disabled by default to avoid affecting timestamp-sensitive logic.

You can enable or disable unified time zone using the following script in the ECS parcel:

```
bash /opt/cloudera/parcels/ECS/k8tz-webhook/configure-k8tz-webhook.sh -h
```

This script modifies the k8tz webhook settings.

Syntax:

```
configure-k8tz-webhook.sh [-i|-h]
```

Options:

- **i** – This option enables the unified time zone feature
- No options – To disable the unified time zone feature, run the configure-k8tz-webhook.sh script without any options.
- Use the **-h** flag to print Help information

To complete the process of enabling the unified time zone feature:

- Restart the workload pods where you want the Cloudera Manager Server timezone to be applied.

-OR-

- Initiate an ECS cluster rolling restart. This will inject the time zone information into all workload pods.

When the unified time zone feature is disabled, all running pods are not affected. To apply the new disabled setting so they run with the default UTC time zone, a pod restart or a rolling restart is required.

Adjusting the expiration time of ECS cluster certificates

The RKE Kubernetes, Vault, and ECS webhook certificate expiration times are set to one year by default. To avoid certificate expiration errors, you may want to extend the expiration times.

About this task



Note:

This topic only applies to internal certificates within ECS. It does not apply to the ingress controller certificate.

- These steps describe how to adjust the expiration time of internal cluster certificates in an existing ECS cluster.
- For a new cluster, if the nodes have been added to Cloudera Manager before creating the ECS cluster, you can edit the `cluster_signing_duration` configuration property in Cloudera Manager before creating the ECS cluster.

Adjusting the expiration time of the RKE Kubernetes cluster certificate

1. In Cloudera Manager, select the ECS cluster, then click ECS.
2. Click the Configuration tab, then use the Search box to locate the `cluster_signing_duration` configuration property.

- The `cluster_signing_duration` configuration property sets the expiration time for the RKE Kubernetes, Vault, and ECS webhook certificates, and is set to 1 year (365 days) by default. In the example below, the certificate expiration has been reset to 5 years (1825 days):

The screenshot shows the Cloudera Manager interface for an ECS cluster (ID: 153-b278). The left sidebar contains navigation options: Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, Data Services (New), Parcels, Running Commands, Support, and a user profile for 'admin'. The main panel displays the 'Configuration' tab for the ECS cluster. A search bar at the top shows 'cluster_signing_duration'. Below the search bar, there are filters for SCOPE, CATEGORY, and STATUS. The SCOPE filter shows 'ECS (Service-Wide)' selected. The CATEGORY filter shows 'Main' selected. The STATUS filter shows 'Non-Default' selected. The configuration table shows the 'cluster_signing_duration' property set to '1825' under the 'ECS (Service-Wide)' scope. A 'Save Changes(CTRL+S)' button is at the bottom right.

- Click Save Changes.
- On the ECS Cluster landing page, click Actions > Refresh Cluster.
- After the Refresh is complete, click Actions > Rolling Restart.
- After the restart is complete, the certificate expiration time is reset to the new value. You can also use the CLI to verify the new certificate expiration setting:

```
[root@host-1 ~]# cat /proc/47803/environ
CDH_PIG_HOME=/usr/lib/pigLD_LIBRARY_PATH=:/opt/cloudera/cm-agent/libCMF
_AGENT_ARGS=CDH_KAFKA_HOME=/usr/lib/kafka
CONF_DIR=/var/run/cloudera-scm-agent/process/1546342871-ecs-ECS_SERVERCDH_
PARQUET_HOME=/usr/lib/parquet
PARCELS_ROOT=/opt/cloudera/parcelsPARCEL_DIRNAMES=ECS-1.5.2-b866-ecs-1.5.2
-b866.p0.46395126LANG=en_US.UTF-8
CDH_HADOOP_BIN=/usr/bin/hadoopCDH_KMS_HOME=/usr/lib/hadoop-kmsCGROUP_GROUP
_CPU=CMF_PACKAGE_DIR=/opt/cloudera/cm-agent/service
ORACLE_HOME=/usr/share/oracle/instantclientMGMT_HOME=/opt/cloudera/cmINV
OCATION_ID=04c94a229a2b4684a95f8ec63783c81e
JSVC_HOME=/usr/libexec/bigtop-utilsCDH_IMPALA_HOME=/usr/lib/impalaKRB5_C
ONFIG=/etc/krb5.conf
CDH_YARN_HOME=/usr/lib/hadoop-yarnCLOUDERA_POSTGRESQL_JDBC_JAR=/opt/clo
udera/cm/lib/postgresql-42.5.1.jar
CDH_SOLR_HOME=/usr/lib/solrHIVE_DEFAULT_XML=/etc/hive/conf.dist/hive-defa
ult.xml
CLOUDERA_ORACLE_CONNECTOR_JAR=/usr/share/java/oracle-connector-java.jarC
GROUP_GROUP_BLKIO=system.slice/cloudera-scm-agent.service
```



```
CGROUP_ROOT_BLKIO=/sys/fs/cgroup/blkioCGROUP_ROOT_CPU=/sys/fs/cgroup/cpu,c
puacctKEYTRUSTEE_KP_HOME=/usr/share/keytrustee-keyprovider
CLOUDERA_MYSQL_CONNECTOR_JAR=/usr/share/java/mysql-connector-java.jarCMF_
SERVER_ROOT=/opt/cloudera/cm
CGROUP_ROOT_CPUACCT=/sys/fs/cgroup/cpu,cpuacctCDH_FLUME_HOME=/usr/lib/f
lume-ng
CATTLE_NEW_SIGNED_CERT_EXPIRATION_DAYS=1825
<snip!>
```

```
[root@host-1 ~]# openssl x509 -in /var/lib/rancher/rke2/agent/serving-kubele
t.crt -noout -text
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 4005696761303552502 (0x379717fb376e51f6)
        Signature Algorithm: ecdsa-with-SHA256
        Issuer: CN = rke2-server-ca@1697759349
        Validity
            Not Before: Oct 19 23:49:09 2023 GMT
            Not After : Oct 17 23:49:10 2028 GMT
        Subject: CN = host-1.rke-1019.kcloud.cloudera.com
        Subject Public Key Info:
            Public Key Algorithm: id-ecPublicKey
            Public-Key: (256 bit)
            pub:
                04:92:81:74:b8:fb:aa:6c:c5:9a:40:2c:5f:91:60:
                35:16:9a:d5:41:b2:bf:d8:29:f4:ed:68:ed:cd:3d:
                87:0e:59:db:27:26:c5:d8:a7:79:c7:23:8f:0b:71:
                c2:f5:d4:36:fe:97:a9:b5:62:ee:9d:9b:6d:ed:25:
                60:fd:26:3a:08
            ASN1 OID: prime256v1
            NIST CURVE: P-256
        X509v3 extensions:
            X509v3 Key Usage: critical
                Digital Signature, Key Encipherment
            X509v3 Extended Key Usage:
                TLS Web Server Authentication
            X509v3 Authority Key Identifier:
                keyid:26:8F:9F:A1:04:CE:2D:04:3A:03:11:87:9D:DF:5A:B7:5C:0
6:72:32
            X509v3 Subject Alternative Name:
                DNS:host-1.rke-1019.kcloud.cloudera.com, DNS:localhost, IP
Address:127.0.0.1, IP Address:10.17.130.15
        Signature Algorithm: ecdsa-with-SHA256
            30:46:02:21:00:fc:5c:89:ab:99:a6:79:33:a9:28:da:a8:47:
            52:cf:1f:43:13:8c:06:2e:23:67:4c:b4:b0:d6:e3:f9:b6:ad:
            50:02:21:00:c7:64:aa:86:97:5a:f3:12:7e:3f:a2:f1:ab:93:
            17:6c:3a:37:34:01:ef:ba:7f:08:85:70:2c:c9:40:e0:30:f5
```

Adjusting the expiration time of the Vault certificate

1. In Cloudera Manager, select the ECS cluster, then click ECS.
2. Click the Configuration tab, then use the Search box to locate the cluster_signing_duration configuration property.

- The `cluster_signing_duration` configuration property sets the expiration time for the RKE Kubernetes, Vault, and ECS webhook certificates, and is set to 1 year (365 days) by default. In the example below, the certificate expiration has been reset to 5 years (1825 days):

The screenshot shows the Cloudera Manager interface for an ECS cluster (ID: 153-b278). The left sidebar contains navigation options: Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, Data Services (New), Parcels, Running Commands, Support, and a user profile for 'admin'. The main panel is titled 'ECS' and shows the 'Configuration' tab. A search bar at the top of the main panel contains 'cluster_signing_duration'. Below the search bar, there are filters for SCOPE, CATEGORY, and STATUS. The SCOPE filter shows 'ECS (Service-Wide)' selected. The CATEGORY filter shows 'Main' selected. The STATUS filter shows 'Non-Default' selected. The main configuration area shows 'Cluster Signing Duration' set to 'ECS (Service-Wide)' with a value of '1825'. A 'Save Changes(CTRL+S)' button is at the bottom right. The bottom status bar shows '1 Edited Value' and 'Reason for change: Modified Cluster Signing Duration'.

- Click Save Changes.
- Copy the `rotate-vault-cert.sh` file to the ECS master host. Set `JAVA_HOME` if needed.
- Run the following command:


```
./rotate-vault-cert.sh APP_DOMAIN
```
- Unseal Vault.
- Restart all of the pods in the CDP namespace.
- If you are using a default self-signed ingress controller certificate, update the ingress controller certificate (follow the steps in the script output).
- You can use the CLI to verify the new certificate expiration setting:

```
root      49076    48970    2 16:49 ?        00:00:10 kube-controller-mana
ger
--flex-volume-plugin-dir=/var/lib/kubelet/volumeplugins --terminated-pod-
gc-threshold=1000 --permit-port-sharing=true
--allocate-node-cidrs=true --authentication-kubeconfig=/var/lib/rancher/
rke2/server/cred/controller.kubeconfig
--authorization-kubeconfig=/var/lib/rancher/rke2/server/cred/controller.
kubeconfig --bind-address=127.0.0.1
--cluster-cidr=10.42.0.0/16 --cluster-signing-duration=43800h
<snip!>
```

```
[root@host-1 ~]# openssl x509 -in vault.pem -noout -text
Certificate:
```

```

Data:
  Version: 3 (0x2)
  Serial Number:
    db:b7:a7:c3:79:86:4c:54:e8:97:49:bf:99:3d:df:a9
  Signature Algorithm: ecdsa-with-SHA256
  Issuer: CN = rke2-server-ca@1697759349
  Validity
    Not Before: Oct 19 23:46:38 2023 GMT
    Not After : Oct 17 23:46:38 2028 GMT
  Subject: O = system:nodes, CN = "system:node:vault.vault-system.svc
; "

  Subject Public Key Info:
    Public Key Algorithm: rsaEncryption
    RSA Public-Key: (2048 bit)
    Modulus:
      00:94:93:2e:9d:5c:01:5a:95:46:b2:9d:aa:23:c4:
      4e:0f:92:07:7e:0e:3a:21:7d:ef:95:e8:09:d3:88:
      38:ac:e9:9f:c2:36:37:04:56:43:87:3a:6f:34:08:
      09:8f:3f:df:31:79:d6:12:db:78:f6:1c:9b:0e:c2:
      d0:f5:25:50:86:37:d5:ff:f7:a0:82:6f:55:d1:ff:
      03:54:f8:ce:8b:02:87:2d:af:3f:71:f8:c4:a9:f0:
      24:50:7b:07:70:3d:7a:be:9d:41:f0:15:2f:56:c3:
      d3:0d:1a:e1:87:8e:69:89:ff:bf:1b:f2:84:87:6c:
      5e:f9:13:8b:2c:5c:de:64:9e:ae:de:6a:f0:7c:ae:
      d9:01:41:aa:39:00:b3:2d:4f:5c:db:fb:2b:80:31:
      88:b5:40:24:e1:06:08:c4:ad:82:70:a1:9e:4c:3e:
      00:0d:61:d9:1a:5c:c7:11:a7:79:68:66:34:b2:c2:
      e9:63:a8:5d:d1:13:be:e6:f1:8f:03:87:3d:be:eb:
      b7:ce:a5:eb:56:81:37:5b:9d:ce:82:34:15:99:16:
      4c:65:20:d9:df:e6:63:56:c2:49:79:e8:66:ce:c1:
      01:9d:87:a2:ba:02:c0:7c:2b:e5:37:30:c5:23:bd:
      87:a1:c8:2b:a9:49:be:67:31:22:8d:a4:68:f9:bd:
      be:23
    Exponent: 65537 (0x10001)
  X509v3 extensions:
    X509v3 Key Usage: critical
      Digital Signature, Key Encipherment
    X509v3 Extended Key Usage:
      TLS Web Server Authentication
    X509v3 Basic Constraints: critical
      CA:FALSE
    X509v3 Authority Key Identifier:
      keyid:26:8F:9F:A1:04:CE:2D:04:3A:03:11:87:9D:DF:5A:B7:5C:0
6:72:32
    X509v3 Subject Alternative Name:
      DNS:vault, DNS:vault.vault-system, DNS:vault.vault-system.
      svc, DNS:vault.vault-system.svc.cluster.local, DNS:vault.localhost.localdoma
      in, DNS:*.apps.host-1.rke-1019.kcloud.cloudera.com, IP Address:127.0.0.1
    Signature Algorithm: ecdsa-with-SHA256
      30:46:02:21:00:d9:5e:38:fc:31:9b:5a:eb:fc:7d:c2:8f:b3:
      54:5e:28:f0:8f:00:eb:36:65:9f:d3:70:ae:a2:79:77:ee:b5:
      f7:02:21:00:f4:e8:6f:c9:bd:bb:92:9d:63:81:69:55:67:8b:
      8a:f3:a4:5d:c1:67:66:b0:40:ff:22:a6:c3:6f:4f:8e:b2:8e

```

Adjusting the expiration time of the ECS webhook certificate

1. In Cloudera Manager, select the ECS cluster, then click ECS.
2. Click the Configuration tab, then use the Search box to locate the cluster_signing_duration configuration property.

- The `cluster_signing_duration` configuration property sets the expiration time for the RKE Kubernetes, Vault, and ECS webhook certificates, and is set to 1 year (365 days) by default. In the example below, the certificate expiration has been reset to 5 years (1825 days):

The screenshot shows the Cloudera Manager interface for an ECS cluster (ID: 153-b278). The left sidebar contains navigation options: Clusters, Hosts, Diagnostics, Audits, Charts, Replication, Administration, Data Services (New), Parcels, Running Commands, Support, and a user profile for 'admin'. The main panel displays the 'Configuration' tab for the 'cluster_signing_duration' property. The property is set to 'ECS (Service-Wide)' with a value of '1825'. A 'Filters' sidebar on the left shows the following counts:

SCOPE	CATEGORY	STATUS
ECS (Service-Wide) 1	Main 1	Error 0
Ecs Agent 0	Advanced 0	Warning 0
Ecs Server 0	Monitoring 0	Edited 1
	Performance 0	Non-Default 1
	Ports and Addresses 0	Include Overrides 0
	Resource Management 0	
	Security 0	

At the bottom, a 'Reason for change' field contains 'Modified Cluster Signing Duration' and a 'Save Changes(CTRL+S)' button is visible.

- Click Save Changes.
- Copy the `rotate-webhook-cert.sh` file to the ECS master host.
- Run the following command:


```
./rotate-webhook-cert.sh APP_DOMAIN
```
- Check for any pods in the Pending state whose status shows that they cannot tolerate the `node-role.kubernetes.io/control-plane` toleration. Restart those pods.
- You can use the CLI to verify the new certificate expiration setting:

```
root      49076    48970    2 16:49 ?        00:00:10 kube-controller-manage
--flex-volume-plugin-dir=/var/lib/kubelet/volumeplugins --terminated-pod-gc-threshold=1000 --permit-port-sharing=true
--allocate-node-cidrs=true --authentication-kubeconfig=/var/lib/rancher/rke2/server/cred/controller.kubeconfig
--authorization-kubeconfig=/var/lib/rancher/rke2/server/cred/controller.kubeconfig --bind-address=127.0.0.1
--cluster-cidr=10.42.0.0/16 --cluster-signing-duration=43800h
<snip!>
```

```
[root@host-1 ~]# openssl x509 -in ecs-tolerations-webhook-cert.pem -noout -text
Certificate:
    Data:
```

```

Version: 3 (0x2)
Serial Number:
    a5:31:94:f4:84:bb:3b:a2:a4:63:8d:ec:de:b5:37:53
Signature Algorithm: ecdsa-with-SHA256
Issuer: CN = rke2-server-ca@1697759349
Validity
    Not Before: Oct 19 23:45:48 2023 GMT
    Not After : Oct 17 23:45:48 2028 GMT
Subject: O = system:nodes, CN = "system:node:ecs-tolerations-webhook
.ecs-webhooks.svc;"
Subject Public Key Info:
    Public Key Algorithm: rsaEncryption
    RSA Public-Key: (2048 bit)
    Modulus:
        00:cc:12:e1:54:b8:aa:42:94:aa:11:a5:f7:35:0e:
        0c:de:76:5b:d5:c6:c1:34:0b:b8:b7:2b:15:08:1d:
        02:44:0f:2e:e1:17:dc:73:6a:e4:6c:df:5b:ac:43:
        97:2e:34:73:f7:c9:6f:cf:c2:a8:52:79:b1:89:ea:
        51:22:e1:41:b8:6a:ba:fd:22:a2:bf:a2:46:a4:8e:
        f5:c6:2d:05:c3:a5:1d:6b:60:da:e8:40:a5:e1:e1:
        5a:55:0e:94:2d:91:dd:71:d1:e9:aa:27:5d:e6:fc:
        ea:5f:ea:c6:8e:52:71:27:ce:c2:a7:1b:10:ca:db:
        db:27:c8:46:6d:14:d1:d0:b3:f5:ab:74:a9:63:8b:
        71:83:31:eb:ad:87:1b:3b:8d:ff:ce:d0:7f:d1:1b:

```

Configuring multiple Base clusters with one ECS cluster

You can configure one Embedded Container Service (ECS) cluster to work with multiple CDP Private Cloud Base clusters managed by separate instances of Cloudera Manager. In order to do this you must first create a combined truststore .pem file that contains the ECS Control Plane truststore .pem file appended with the certificate files of each of the CDP Private Cloud Base clusters.

About this task

Use the following steps to configure one ECS cluster to work with multiple CDP Private Cloud Base clusters:

1. Append the ECS Control Plane truststore .pem file with the certificate files from the additional CDP Private Cloud Base clusters.
2. Register an ECS environment with each of the additional CDP Private Cloud Base clusters.
3. Create data services within each environment.

Step 1: Append the ECS Control Plane truststore .pem file with the certificate files from the Base clusters

1. On the ECS Control Plane, run the following kubectl command to get the contents of the configmap:

```
kubectl get configmap cdp-private-installer-truststore -n cdp -o yaml >
cdp-private-installer-truststore.yaml
```

2. Copy the truststorePEM content, decode it, and store it in a file. For example:

```
echo LS0tLS1CRUdJTiBDRVJUSUZJQ0FURS0tLS0tCk1JSURhakNDQWxJQ0NRRG5iNnhmK0d
QR1l6QU5CZ2txaGtpRz13MEJBUXNGQURCWk1Rc3dDUV
lEVlFR0V3S1YKVXpFTTE1Ba0dBWVVFQ0F3Q1EwRXhDekFKQmdOVk1JBY01BbE5ETVEwd0N3WU
RWUVFLREFSRFRFULNNUXd3Q2dZRApWUVFMREFOQ1RGS
XhFekFSQmdOVk1JBTU1DaW91YUhkNEUxTnBkR1V3SGhjTk1qTXhNVEV3TVRRME1qUXdXaGNOC
kl1qUXhNVEE1TVRRME1qUXdXakFwTVJNd0VRWURWUVFE
REFvcUxtaDNlQzV6YVhSbE1JSUJvakFOQmdrcWhraUcKQXcwQkFRRUZBQU9DQV4k4QU1JSUJp
Z0tDQVlFQS9lZk1kK05lQ0tWUTF1M05qK3Z0RGFRV0p
```

```
JcUhFbVcxOf1pYgpbQUdiYmlvYi9YyN0aTRINU81MXV3SjJ1cWowaktUM3dBU3l0UG0yS0p
1RE9vVXMveWhJc0xuK3VOWlMzd292CkNxSk5RcWpRT3
N2RUVITU5ZZ3JOWExMclhlbHZHTXl4aG16bVf1SEhHTkZhclDENVkwD1laMVVlaG00a0pUUT
UKTFhoZmlJVjJlTUJieE4ySVB2WU1TVlAvYmo4ekF3a
k50OHQvVUhhafRTEw1jUktEWitsMGxoeGt0cHppzdmxmcQo4eXNcVTBBQ2MvbWp2bGNWS0xyN
VVRSTRadVNFb2ZRKlQyaEpITEZnQ0N4bFJvcWN5aFo0
QmtlZmZwaUhIOGJHCm9kd2tSaHRRMVfJcFFxSk1CLytCOWNZbkFjYlBFaHlXekh1TGlqakl5
VTZOYWZ3SmpoTG1SVmptRmpWNzNvZmgKanJ4V1BtVyt
FSDJZODRWK3RpOVDiZE5LQW9KNzU4bzZaSmJsc3ZBRVBNVytBVmw2clFMTTFPZXNlUTNtczc
xMwpWOENObFBWVEQ0UGdpaythOG1YV3FWZkVZN2F1V3
N1YnIwUkIyeFlIWHBhd2lWdWxrSjdYRURHOEpmN2hFNzRqCkRhMlJaeWN5YXdScGF3SXV2V1
kwWGtoSktOOTNBZ0lCQUFFd0RRWUpLb1pJaHJzjTfFRR
UxCUUFZE2dFQkFDcTcKSU5R2lnKy9iUVB3enhmUmF6dlhXM09mT3M1UjNnU0hGeDRmS1BXV
lN5TjEwaW50bmdxejd4R2dYVnBpRDdWNApQRGVXZFRZ
MjdHN2w3ZHBjek1FS2ptN25XOU3RW05S3dyRndWRWh0OWeZnjVvUnhqTzA3Y09VanZYaEwy
dkx1Cnk1eHRYZlJyZXlPalNmZDVxcnlKVlBoMDBHB0N
UWTViMy9wK25saWJUUmNkY29mQkFTU0VhbnhaVDJoc1B2V3kKSG9PVkVGSmlrTnVxRHJhS2Y
ySlFxrRn4aGs0MFivUW9LVUpKUTgzUWixZHBmWWVCDe
9lWXRvNExmQWV3Y0RuRwpFWUQvYVp1blglwU2cxRTRoRS9NaUNFN2R6ZzY4TVVPeWVBV1pCel
JuMHBEZlVtanpTOUNndi9GQ240MjV0QnR5Cis5anYlW
it3TVNkd1VZL2VudEE9Ci0tLS0tRU5EIEF5F1RJRklDQVRFLS0tLS0KLS0tLS1CRUDJTiBDR
VJUSUZJQ0FURS0tLS0tCk1JSURlekNDQWlPZ0F3SUJB
Z0lVQWRide1lQ3JycVRMYlUzRzhPakZRUW5YNGY4d0RRWUpLb1pJaHJzjTfFRUwKQlFBd1dU
RUxNQWtHQTFVRUJoTUNWVkl4Q3pBSk1JN1ZlZCQWdNQWt
OQk1Rc3dDUVlEVlFRSERBSlRRekVOTUFzRwpBMVVFQ2d3RVEweEVVakVNTUFvR0ExVUVDd3d
EUWt4U01STXdfUUVlEVlFRERBb3FMbWgzZUM1emFYUm
xNQjRYCkRUSXpNVEV4TURFek1UTXpOVm9YRFRJMU1URXdPVEV6TVRNek5Wb3dXVEVMTUFrR0
ExVUVCaE1DVlZNeEN6QUoKQmdOVk1JZ01Ba05CTVFzd
0NRWURWUWFIREFKVFF6RU5NQXNHQTFVRUNnd0VRMHhFVWpFTU1Bb0dBMVVFQ3d3RApRa3hTT
VJND0VRWURWUWFEREfvcUxtaDNlQzV6YVhSbE1JSUJj
akFOQmdrcWhraUc5dzBCQVFFRkFBT0NBUThBCK1JSUJJDZ0tDQVFFQXczQXBYeXg4dkxXSzVZq
SlpLZzNpb29XcGdtNjZwN2gxWCtRWUUVVZ0Q0VEc3dkZ
2OGNUckKdzlaZlVpcW1zUTVJRlZxRk5lcEFpSFBteUxscDl1d1RhTEThdM9IZ2pXU0p1K2d
waUdiMHJiR1hkM3ltYkw5Rwp2Sm1pNmtPZW9SeHpQbk
N5SVVEa3NmU3kzde5pWlNRRFRubmhUWk9Zc2tmbDdZK1VYaVJVS2NBNEExkWTBwSTVJCnmpRl
R0cW5qM0o4SnJ6d0dJd1NoK0ZNdHRYWFQ5WFI5bzVpL
0M2cWh0L1JwbEx3QTB6ZVlYSdhkNj12Ykw4T1EKemREeXZlcmptRXZjs3F1bGo4NU1CSTZwc
VRGb2lQcEp5VVlxS0cwN2U1WDN0QmZiVzk2QXdxYT1BT
SFd0QlplndwpyeTVFbZRxWVRJMGZmYlFCS3ZIVElZYTd3T0xmRzAvK3J3SURBUUFcb3pzd09U
QUxCZ05WSFE4RUJBTUNCREF3CkV3WURWUjBsQkF3d0N
nWU1Ld1lCQlFVSEF3RXdGUUVlEVlIwUk1JBNHdESU1LS2k1b2QzZ3VjMmwwWlRBTk1Jna3EKAgt
pRz13MEJBUXNGQUFPQ0FRRUFTkZfZU1g5M2k1Q1FPQ1
FIVVZ2Y2M1OWFmb2Y3SnJxcGNAN0NOaGJXMzc4Zgo3RTNpTjhBY1BNQ0dvZllTeWFrblQxVl
kwDDNiVXhtSTFSdXdeUXNDU3U1MmlhYnhIVUhrOFBEQ
jk5NTRxL3RtCkh4MXpVR0VURkZaZhdKb0dDMk14Ui9WdU9wbExza2hEc0ZJZmpaZC81clVrL
lQvMUxUaC8zMExBbGhPVzNtek8KZFJWWC9LR2QyWGZ3
SFNZq3FRTFk4WGZQM0d3WHgrTmVUY09vTEQycXYvYw1kMnY1dlVtdXpONzErZjR3bXVvbWpa
Z1JiYk9OSkMvdzVzV3MvWVRaODd1M1JNUWExd2gvcKl
YMk1QMzNTMG1SeHJkSXlpeGMxamF6ZTYxWmRUUnk5Ck9NQ2RmZEpGNFE1RndmODdWSWpYZXdx
PemdQVnFJVGvNVW1vcy9HR0p0UT09Ci0tLS0tRU5EIE
NFU1RJRklDQVRFLS0tLS0= | base64 -d > cdp-private-installer-truststore.pem
```

3. Obtain the truststore .pem file from the first additional Cloudera Manager host from /var/lib/cloudera-scm-agent/agent-cert/cm-auto-global_cacerts.pem or /opt/cloudera/CMCA/trust-store/cm-auto-global_cacerts.pem and copy the contents.
4. Append the cdp-private-installer-truststore.pem file created previously with the contents of the Cloudera Manager .pem file.
5. Repeat the previous two steps for all additional Cloudera Manager hosts you would like to register environments with.

6. Log in to the ECS cluster Management Console and click Administration > CA Certificates. Select Datalake in the CA Certificate Type drop-down, click Choose File, then select the appended cdp-private-installer-truststore.pem file and click Upload. Click Save to save your changes.

You can also use the following CLI commands to upload the cdp-private-installer-truststore.pem file and update the global truststore with the encoded certificate file content:

```
cat cdp-private-installer-truststore.pem | base64
cdp environments --set-environment-setting --settings truststorePEM=<base64 encoded CM cert> --no-verify-tls
```

Step 2: Register an ECS environment with each of the additional Base clusters

1. Log in to the ECS cluster Management Console and [Register an environment](#) for the first additional Base cluster using the applicable Cloudera Manager URL and credentials.
2. Repeat the previous step for the rest of the additional Base clusters.

Step 3: Create data services within each environment

Refer to the following topics to create the data services of your choice in each environment:

- [Adding a Cloudera Data Engineering service](#)
- [Activate ECS environments \(CDW\)](#)
- [Provision an ML Workspace](#)

GPU node labeling on ECS

You can use NVIDIA Feature Discovery to generate labels for the set of GPUs available on ECS nodes. You can use these node labels to assign workloads to specific GPU devices. This feature is enabled by default on ECS.

Using GPU node labeling on ECS

Information about using GPU node labeling is available on the [NVIDIA GPU feature discovery](#) page.

Known Issues and Limitations

- GPU node labeling is only supported for GPU cards manufactured by NVIDIA.
- If an ECS node has [multiple GPUs](#), not all of the GPUs will be labeled. The last GPU as per lspci will be labeled.
- If an ECS node is provisioned from a provisioner with virtual GPUs (AWS, Azure, etc.) the nodes will not be labeled with the GPU information.