

Cloudera Management Console Release Notes

Date published: 2019-08-22

Date modified: 2025-08-18



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

What's new..... 4

 August 18, 2025..... 4

 July 17, 2025..... 4

 June 4, 2025..... 4

 May 29, 2025..... 5

 May 7, 2025..... 5

 April 2, 2025..... 5

 March 3, 2025..... 6

 February 21, 2025..... 6

 January 30, 2025..... 7

 January 15, 2025..... 7

 Older releases..... 7

 2024..... 7

 2023..... 13

 2022..... 17

 2021..... 24

 2020..... 37

 2019..... 41

Kubernetes Ingress NGINX Controller vulnerabilities..... 42

Log4j vulnerabilities..... 43

Image catalog updates..... 43

Known issues for Cloudera Management Console..... 44

Behavioral changes in Cloudera Management Console..... 50

What's new

This section lists major features and updates for the Cloudera Management Console service.

August 18, 2025

This release of the Cloudera Management Console service introduces the following changes:

Transitioning from Azure default outbound access

As previously communicated in [Cloudera Customer Advisory-866](#), Microsoft is planning to retire the default outbound access connectivity for virtual machines in Azure on 30 September, 2025. You must transition your Cloudera environment on Azure to use an explicit method for outbound internet connectivity while creating any new virtual machines. This ensures enhanced security, predictability, and reliability for your clusters.

For more information, see the [Transitioning from Azure default outbound access](#) documentation.

July 17, 2025

This release of the Cloudera Management Console service introduces the following changes:

Support for Spain Central Azure region

The Spain Central Azure regions is now supported. You can register Azure environments and provision Cloudera Data Hub clusters in these regions. See updated [Supported Azure regions](#).

Secure binds for LDAP in FreeIPA

By default, secure authentication is available for LDAP (port 636) on FreeIPA instances for new and existing environments, which means that the secure authentication is automatically used for FreeIPA, Data Lake and Cloudera Data Hub within a Cloudera environment.

For more information, see the [Secure binds for LDAP in FreeIPA](#) documentation.

June 4, 2025

This release of the Cloudera Management Console service introduces the following changes:

Cluster Connectivity Manager v1 - End of Life (EoL) update

As previously communicated in [Cloudera Customer Advisory-824](#), Cluster Connectivity Manager v1 has reached its EoL by April 10, 2025 in favor of Cluster Connectivity Manager v2.

The Cluster Connectivity Manager v1 infrastructure has been shut down on April 10, 2025. All traffic on the Cluster Connectivity Manager v1 channel has been blocked. Starting June 4th, Cloudera has initiated a phased decommissioning of the infrastructure to remove Cluster Connectivity Manager v1 from Cloudera on cloud and on premises.

For more information about the discontinuation of Cluster Connectivity Manager v1, see [Cloudera Customer Advisory-824: Cluster Connectivity Manager v1 will be discontinued on April 10, 2025](#).

May 29, 2025

This release of the Cloudera Management Console service introduces the following changes:

Secret rotation - GA

To add extra measures of security, you can rotate secrets, like database passwords or FreeIPA admin password using CLI commands. These secrets are managed and created by the Cloudera Control Plane. By using the following commands the secrets can be rotated to achieve more secure deployments.

For more information, see the following documentations:

- [Rotating FreeIPA secrets](#)
- [Rotating Data Lake secrets](#)

Restarting instances from Cloudera Management Console

You can now restart Data Lake instances using the Cloudera Management Console. For more information, see the [Restarting Data Lake instances](#) documentation.

Updated Minimal policy for AWS

The Minimal policy that is required for accessing AWS resources for Cloudera on cloud has been updated to limit the permissions for deleting resources. The updated Minimal policy can be found [here](#) or in the [Reduced access policy definition](#) documentation.

Updated Minimal permissions for GCP

The Minimal permissions for GCP that are required for Cloudera environments, Data Lake, FreeIPA and Cloudera Data Hub are updated with a new set of permissions. For more information, see the [Permissions for the provisioning credential's service account](#) documentation.

May 7, 2025

This release of the Cloudera Management Console service introduces the following changes:

Azure certificate credentials and managed identity for AKS

You can use certificate based authentication with managed identity when using AKS clusters for Data Services. Certificate based authentication uses asymmetric authentication based on private-keys and public-keys, and does not rely on sharing secrets. In order to use certificate credentials, you need to create user-assigned managed identity. The managed identity can be attached to the Cloudera environment.

For more information see the following documentations:

- [Prerequisites for the provisioning credential](#)
- [Generate and register Azure certificate credentials](#)
- [Using a user-managed identity for AKS clusters](#)
- [Encrypting VM disks with customer managed keys](#)

April 2, 2025

This release of the Cloudera Management Console service introduces the following changes:

Creating new network for AWS environments is removed

The option to create a new network when registering a Cloudera environment on AWS is no longer available. To register a Cloudera environment on AWS, you can use your existing VPC and subnets already available in AWS.

This change impacts the following documentations:

- [Register an AWS environment from the Cloudera UI](#)
- [AWS Requirements - VPC and subnets](#)
- [Determining the CIDR range](#)

For quickly deploying Cloudera environments on AWS, see the [Deploy Cloudera using Terraform](#) documentation.

Secret rotation (Preview)

To add extra measures of security, you can rotate secrets, like database passwords or FreeIPA admin password using CLI commands. These secrets are managed and created by the Cloudera Control Plane. By using the following commands the secrets can be rotated to achieve more secure deployments.



Note: The documentation of Secret rotation has been moved to the following locations as the feature is Generally Available (GA):

- [Rotating FreeIPA secrets](#)
- [Rotating Data Lake secrets](#)
- [Rotating Cloudera Data Hub secrets](#)

March 3, 2025

This release of the Cloudera Management Console service introduces the following changes:

Added Stale cluster status

Stale cluster status has been added to the list of cluster statuses to express a status where we have no information from the cluster and it has been unreachable for more than 30 days and that the Data Lake's status is outdated or no longer reflects its current state. For more information, see [Data Lake status options](#).

Added get-operation command to obtain operation status

A new command, get-operation is now available to get the status of the latest operation or a specified earlier operation performed on an environment or a Data Lake cluster. For more information, see [Monitoring a Data Lake](#), [Monitoring an AWS environment](#), [Monitoring an Azure environment](#), and [Monitoring a GCP environment](#).

Support for Italy North, New Zealand North, and Poland Central Azure regions

The Italy North, New Zealand North, and Poland Central Azure regions are now supported. You can register Azure environments and provision Cloudera Data Hub clusters in these regions. See updated [Supported Azure regions](#).

February 21, 2025

This release of the Cloudera Management Console service introduces the following changes:

Slack integration for Notifications - Technical Preview

Slack is added as a communication channel for the Notification service beside in-app and email. After adding the Cloudera Notifications application in Slack, you can receive the resource notifications as slack messages. For more information, see the [Setting up Slack for resource notifications](#) documentation.

January 30, 2025

This release of the Cloudera Management Console service introduces the following changes:

Imported Compute Engine images encrypted with CMEK

If you set a CMEK for your GCP environment, then the imported Compute Engine images will be encrypted with the CMEK instead of the default Google-managed key. For more information, see [Adding a customer managed encryption key for GCP](#).

Enabling Secure Boot option for GCP (Preview)

VMs on GCP are created without the Secure Boot option enabled by default. You can request to have the Secure Boot option for subsequently provisioned VMs. For more information, see [Enabling Secure Boot option for GCP](#).



Note:

You need to contact Cloudera Support to have this feature enabled.

January 15, 2025

This release of the Cloudera Management Console service introduces the following changes:

Added support for upgrading CMK enabled Azure Single Servers to Flexible Server

A managed identity is required for being able to upgrade a CMK enabled Azure Single Server to Azure Flexible Server. It is now possible to add a managed identity to an environment with Azure Single Server that is already encrypted with a Customer Managed Key thus making the upgrade from Azure Single Server to Azure Flexible Server possible for environments already configured with a CMK.

For more information, see [Upgrading Azure Single Server to Flexible Server Prerequisites](#).

Older releases

Overview of new features, enhancements, and changed behavior introduced in earlier releases of Management Console.

2024

December 17, 2024

This release of the Cloudera Management Console service introduces the following changes:

Deploying Cloudera in multiple GCP availability zones (Preview)

You can optionally choose to deploy Data Lake, FreeIPA, and Cloudera Data Hub clusters across multiple availability zones (multi-AZ). With multi-AZ support, newly created GCP environments, enterprise Data Lakes and Cloudera Data Hub clusters using HA templates can be deployed across multiple availability zones of the selected GCP region. This provides fault tolerance during the extreme event of an availability zone outage. For more information, see [Deploying Cloudera in multiple GCP availability zones](#).



Note:

You need to contact Cloudera Support to have this feature enabled.

Added pd-extreme and pd-balanced as supported GCP disktypes

The following disk types have been added to supported GCP block storage types:

- pd-extreme
- pd-balanced

You can review all supported disk types for GCP at [Supported GCP block storage](#).

Added EnvironmentPrivilegedUser to environment resource roles

The new EnvironmentPrivilegedUser resource role grants permission to execute privileged operating system actions on Data Lake, FreeIPA, and Cloudera Data Hub virtual machines.

For more information, see [Resource roles](#).

December 10, 2024

This release of the Cloudera Management Console service introduces the following changes:

Cloudera Runtime 7.3.1

Cloudera Runtime 7.3.1 is now available and can be used for registering an environment with a 7.3.1 Data Lake and creating Cloudera Data Hub clusters. For more information about the new Cloudera Runtime version, see [Cloudera Runtime](#). If you need to upgrade your existing Cloudera environment, your upgrade path may be complex. To determine your upgrade path, refer to [Upgrading to Runtime 7.3.1](#) documentation.

December 9, 2024

This release of the Cloudera Management Console service introduces the following changes:

Compute Cluster enabled environments

The following changes have been introduced to compute cluster enabled environments:

- Suspending and resuming Compute Clusters in Cloudera environments on AWS

For more information, see [Suspending and resuming Compute Clusters](#).

October 29, 2024

This release of the Cloudera Management Console service introduces the following changes:

Azure Single Server to Azure Flexible Server upgrade

Single Server on Microsoft Azure databases used by Data Lakes and Cloudera Data Hub clusters can now be upgraded to Azure Flexible Server. During the upgrade process from PostgreSQL version 11 to PostgreSQL 14, Azure Single Server will be upgraded to Azure Flexible Server. For more information, see [Upgrading Azure Single Server to Flexible Server](#).

Compute Cluster enabled environments

The following changes have been introduced to compute cluster enabled environments:

- You can provide the **Worker Node Subnets** for compute cluster enabled environments on AWS and Azure
- You can provide the **AKS Private DNS Zone ID** for compute cluster enabled environments on Azure

For more information, see [Using compute cluster enabled environments on AWS](#) and [Using compute cluster enabled environments on Azure](#).

New configuration property for non-transparent proxy

Inbound Proxy CIDR has been introduced for configuring non-transparent proxy in Cloudera environments to allow communication with the Kubernetes server when defining the proxy with FQDNs.

For more information, see [Using a non-transparent proxy](#).

Receiving resource notifications

Notifications include automatically generated service and resource related alerts, such as cluster state changes and events, upgrade alerts, resource exhaustion and consumption notifications. Notifications can be received by users of a tenant who have subscribed to the resource events of a Cloudera service.

For more information, see the [Receiving notifications](#) documentation.

October 9, 2024

This release of the Cloudera Management Console service introduces the following changes:

Support for Sweden Central region

The Sweden Central Azure region is now supported. You can register Azure environments and provision Cloudera Data Hub clusters in that region. See updated [Supported Azure regions](#).

September 26, 2024

This release of the Cloudera Management Console service introduces the following changes:

Receiving announcements

You can subscribe to receive announcements and notifications in Cloudera on cloud about various events from product updates to data service specific alerts. Announcements include product related announcements that can include updates related to End of Life (EOL), End of Support (EOS), Technical Service Bulletins (TSBs), and maintenance updates.

For more information, see the [Receiving notifications](#) documentation.

Compute Cluster enabled environments

Compute Clusters enable you to deploy a containerized platform on Kubernetes for Data Services and shared services. The Compute Cluster architecture offers simplified management, enhanced efficiency, and centralized control that leads to faster deployments, reduced configuration errors and improved system reliability. As multiple Data Services can optionally share the same Compute Cluster, it also lowers the cost of ownership.

For more information, see [Using compute cluster enabled environments on AWS](#) and [Using compute cluster enabled environments on Azure](#).

September 3, 2024

This release of the Cloudera Management Console service introduces the following changes:

Support for Private Link for Azure Flexible Server

Azure Database for PostgreSQL Flexible Server allows a highly available database to be deployed for Data Lake and Cloudera Data Hub clusters. It is no longer required to use delegated subnets for Flexible Servers to be used with private access as Private Link for Azure Flexible Server is now supported in Cloudera.

For more information, see [Using Azure Database for PostgreSQL Flexible Server](#) and [Private setup for Azure Flexible Server](#).

Rebuilding FreeIPA (Preview)

You can now rebuild FreeIPA in case all the instances are lost or the LDAP database is damaged beyond repair. For more information



Note:

You need to contact Cloudera Support to have this feature enabled.

For more information, see [Rebuilding FreeIPA](#).

June 12, 2024

This release of the Cloudera Management Console service introduces the following changes:

Support for Hyderabad and Calgary AWS regions

The Asia Pacific Hyderabad AWS region and the Calgary AWS region is now supported. You can register Cloudera environments and provision Cloudera Data Hub clusters in that region. See updated [Supported AWS regions](#).

Support for Madrid Google Cloud region

The Madrid Google Cloud region is now supported. You can register Cloudera environments and provision Cloudera Data Hub clusters in that region. See updated [Supported GCP regions](#).

May 15, 2024

This release of the Cloudera Management Console service introduces the following changes:

Updating instance metadata to IMDSv2

CDP now uses IMDSv2 for accessing EC2 instance metadata on all newly created Data Lakes, FreeIPA clusters, and Cloudera Data Hub clusters. Previously created clusters using IMDSv1 can now be updated to IMDSv2. For more information, see [Updating instance metadata to IMDSv2](#).

April 18, 2024

This release of the Cloudera Management Console service introduces the following changes:

Configuring a CMK for data encryption in Azure Database for PostgreSQL Flexible Server

You can optionally use a customer managed encryption key (CMK) for data encryption in the Azure Database for PostgreSQL Flexible Server database instance used by Cloudera. For more information, see [Configuring a CMK for data encryption in Azure Database for PostgreSQL Flexible Server](#).

Disk vertical scaling in Azure (Preview)

Disk vertical scaling (that is, disk type change and resizing) is now supported by Cloudera for Data Lakes and Cloudera Data Hub clusters running in Azure. Previously, only AWS support was available for this feature in Cloudera. For more information, see [Disk Vertical Scaling - Disk Type Change and Resizing in AWS and Azure](#).



Note: You need to contact Cloudera to have this feature enabled.

April 3, 2024

This release of the Cloudera Management Console service introduces the following changes:

Cloudera Runtime 7.2.18

Cloudera Runtime 7.2.18 is now available and can be used for registering an environment with a 7.2.18 Data Lake and creating Cloudera Data Hub clusters. For more information about the new Cloudera Runtime version, see [Cloudera Runtime](#). If you need to upgrade your existing Cloudera environment, your upgrade path may be complex. To determine your upgrade path, refer to [Upgrading to Runtime 7.2.18](#) documentation.

RHEL replaces CentOS as default OS

As of June 30, 2024, CentOS reaches End of Life (EOL), and consequently, Cloudera Runtime 7.2.18 supports RHEL 8 only. New deployments of Data Lakes and Cloudera Data Hub clusters with Cloudera Runtime 7.2.18 and upgrades to 7.2.18 are only possible with RHEL 8. Data Lake and Cloudera Data Hub clusters running Cloudera Runtime

7.2.17 support both CentOS 7 and RHEL 8. Earlier Cloudera Runtime versions support CentOS 7 only. Cloudera will not publish any updates or fixes for CentOS-based images after June 2024.

As part of the upgrade process of FreeIPA, Data Lake, and Cloudera Data Hub clusters, you have the option to upgrade the operating system (OS) on the virtual machines (VMs) from CentOS 7 to Red Hat Enterprise Linux 8 (RHEL 8). For more information, see [Upgrading from CentOS to RHEL](#).

Discontinuation of Medium Duty Data Lake

Starting with Cloudera Runtime 7.2.18, Medium Duty Data Lake is discontinued and is replaced by Enterprise Data Lake. In order for existing Data Lakes to be upgraded to Runtime 7.2.18, they must be using Enterprise or Light Duty Data Lake.

Enterprise Data Lakes are a redefined version of Medium Duty Data Lakes that still offer failure resilience, but utilize resources and allocate memory more efficiently than a Medium Duty Data Lake at the same cost. Enterprise Data Lakes can handle more intensive workloads than Medium Duty Data Lakes and when deployed in Multi-AZ mode, remain operational during an availability zone outage.

If you are using Medium Duty Data Lake and would like to upgrade to Cloudera Runtime 7.2.18, you will first need to upgrade to 7.2.17 first, and then resize your Data Lake to Enterprise Data Lake. For more information, see [Upgrading from Medium Duty to Enterprise Data Lake](#).

Support for Amazon S3 Express One Zone buckets

Starting with Runtime 7.2.18, CDP supports using S3 Express One Zone buckets for data storage. If you have additional data buckets that you would like to use for Cloudera Data Hub cluster workloads and you do not need zone redundancy, you may use S3 Express buckets, for example for faster processing of temporary data. For more information, see [Using S3 Express One Zone for data storage](#).

Rolling upgrade support for the Data Lake

With the release of Cloudera Runtime 7.2.18, rolling upgrades for certain Data Lakes are now available. Rolling upgrades for the Data Lake are limited to certain Cloudera Runtime versions and shapes of Data Lakes. For more information, see [Rolling upgrades](#).

March 24, 2024

This release of the Cloudera Management Console service introduces the following changes:

IMDSv2 used by default for new clusters on AWS

On AWS, Instance Metadata Service Version 2 (IMDSv2) is used for all newly created Data Lakes, FreeIPA, and Cloudera Data Hub clusters. Previously, IMDSv1 was used.

March 14, 2024

This release of the Cloudera Management Console service introduces the following changes:

RHEL 8 enabled by default

RHEL 8 is now used as a default operating system for all newly created Data Lake, FreeIPA, and Cloudera Data Hub cluster running Cloudera Runtime 7.2.17 and newer.



Note: As of June 30, 2024, CentOS reaches End of Life (EOL), and consequently, Cloudera Runtime 7.2.18 only supports RHEL 8-based images. New deployments of Data Lakes and Cloudera Data Hub clusters with Cloudera Runtime 7.2.18 and upgrades to 7.2.18 are only possible with RHEL 8.

Data Lake and Cloudera Data Hub clusters running Cloudera Runtime 7.2.17 support both CentOS 7 and RHEL 8. Earlier Cloudera Runtime versions supported CentOS only, but Cloudera will not publish any updates or fixes for CentOS-based images after June 2024. Cloudera Runtime 7.2.18 and newer only support RHEL.



Note: On Azure, RHEL 8 images are only available via Azure Marketplace. See [CDP images hosted in Azure Marketplace](#).

The RHEL images, which are CIS Level 1 compliant, are provided with an embedded RHEL license that is restricted for Cloudera on cloud use case. Cloudera's RHEL repository connection is activated during Cloudera deployment and the OS is installed from that repository. Only limited OS packages are hosted in this RHEL repository, but additional OS packages can be requested by filing a support ticket with details on software name, version, and software URL (preferred) or a list of missing OS packages.

March 12, 2024

This release of the Cloudera Management Console service introduces the following changes:

Modify list of AZs available for new Cloudera Data Hub clusters

As an enhancement of the previously released functionality for launching Cloudera environments on Azure in multiple availability zones (AZ), you can now edit an existing environment to modify the availability zones for all newly created Cloudera Data Hub clusters in that environment. See [Deploying CDP in multiple Azure availability zones](#).

Data Lake root volume size was increased to 200 GB

Data Lake's root volume size was increased from 100 GB to 200 GB.

February 6, 2024

This release of the Cloudera Management Console service introduces the following changes:

Tag filtering for Cloudera usage insights

You can now use tags to filter your usage insight based on user-level tags of clusters in your Cloudera environment. For more information, see [CDP credit consumption and usage insights](#).

January 24, 2024

This release of the Cloudera Management Console service introduces the following changes:

Deploying CDP in multiple Azure availability zones

You can now deploy your Cloudera environment, enterprise Data Lake and Cloudera Data Hub clusters on Azure across multiple availability zones. This is an optional configuration that is not enabled by default. For more information, see [Deploying CDP in multiple Azure availability zones](#).

January 5, 2024

This release of the Cloudera Management Console service introduces the following changes:

Azure Database for PostgreSQL Flexible Server

In this release, Cloudera introduces Azure Database for PostgreSQL [Flexible Server](#). New Azure environments automatically use the Flexible Server with public endpoints but as a best practice for production you should configure [Private Flexible Server setup](#).

With the release of this feature, you must add the following permissions on the scope of the single resource group to your custom role:

```
"Microsoft.DBforPostgreSQL/flexibleServers/read",
"Microsoft.DBforPostgreSQL/flexibleServers/write",
"Microsoft.DBforPostgreSQL/flexibleServers/delete",
"Microsoft.DBforPostgreSQL/flexibleServers/start/action",
"Microsoft.DBforPostgreSQL/flexibleServers/stop/action",
"Microsoft.DBforPostgreSQL/flexibleServers/firewallRules/write"
```

For more information, see [Using Azure Database for PostgreSQL Flexible Server](#).

Support for Azure Qatar Central region

Cloudera now supports launching environments and Cloudera Data Hub clusters in Azure Qatar Central region. See [Supported Azure regions](#).

2023

December 20, 2023

This release of the Management Console service introduces the following changes:

VHD images hosted in Azure Marketplace

Cloudera now publishes VHD images on Azure Marketplace for each minor Runtime release (for example, 7.2.17) and CDP uses these images by default during environment and Data Hub creation. In order for CDP to be able to load these images, customers need to accept Azure Marketplace terms and conditions either via CDP web UI (recommended) or Azure CLI. Note that RHEL 8 images on Azure are only available via Azure Marketplace.

In order to use this feature, you need to Grant the service principal additional Azure permissions:

1. On the scope of your Azure subscription:

```
"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/write",  
"Microsoft.MarketplaceOrdering/offerTypes/publishers/offers/plans/agreements/read"
```

2. On the scope of the CDP Azure resource group:

```
"Microsoft.Resources/deployments/whatIf/action"
```

For more information, see [CDP images hosted in Azure Marketplace](#).

December 1, 2023

This release of the Management Console service introduces the following changes:

Additional permissions needed in the cross-account role

You should update the AWS cross-account role to include `ec2:DescribeLaunchTemplateVersions` policy action to allow CDP update AWS Launch Template UserData. CDP Public Cloud cluster operations such as Cluster Connectivity Manager version upgrades, proxy, and Salt credential management require updating the AWS Launch Template UserData. Without this capability, node repair, upgrade, and scale operations performed on the CDP Public Cloud environment, Data Lake (SDX), and Data Hub may fail.

November 1, 2023

This release of the Management Console service introduces the following changes:

Resize an existing Data Lake from single-AZ to multi-AZ

As part of the Data Lake resizing via CDP CLI, you can optionally resize a Data Lake from single-AZ to multi-AZ by adding the `--multi-az` flag to the Data Lake resize command. This is available via CDP CLI only. For more information, see [Data lake scaling](#) and [Scaling the Data Lake through the CDP CLI](#).

October 30, 2023

This release of the Management Console service introduces the following changes:

Qatar Doha (me-central1) and KSA (me-central2) GCP region support

CDP adds support for launching environments, Data Lakes, and Data Hubs in the Qatar Doha (me-central1) and KSA (me-central2) GCP regions. See updated [Supported GCP regions](#).

October 18, 2023

This release of the Management Console service introduces the following changes:

Elastic load balancer deletion protection on AWS

When Data Lakes and Data Hubs are created, CDP deploys load balancers on AWS for endpoint stability. With the release of this update, all newly created load balancers for Data Lakes and Data Hubs on AWS are configured with [deletion protection](#) enabled, provided that your cross-account policy has the required permissions.

Cloudera has updated the AWS cross-account policy definition with the additional `elasticloadbalancing:ModifyLoadBalancerAttributes` permission required to set and remove the load balancer deletion protection flag. If you would like to use this feature, please update your cross account policy on AWS by adding the `elasticloadbalancing:ModifyLoadBalancerAttributes` permission. If you are using the old policy definition, this new feature will not be available (that is, the deletion protection will not be set).

October 9, 2023

This release of the Management Console service introduces the following changes:

Resizing to a Enterprise Data Lake

You can now resize a light or medium duty Data Lake to an Enterprise Data Lake (EDL). Resizing to an EDL is available only on Runtime 7.2.17 and above. For more information, see [Data Lake resizing](#).

October 5, 2023

This release of the Management Console service introduces the following changes:

Public Endpoint Access Gateway for GCP

You can enable Public Endpoint Access Gateway during GCP environment registration. For more information, see [Public Endpoint Access Gateway](#).

September 20, 2023

This release of the Management Console service introduces the following changes:

Support for il-central-1 AWS region

The Tel Aviv (il-central-1) AWS region is now supported. You can register CDP environments and provision Data Hubs in that region. See updated [Supported AWS regions](#).

September 6, 2023

This release of the Management Console service introduces the following changes:

Audit event archiving for GCP

You can now configure audit event archiving for Google Cloud Platform.

For more information, see [GCP setup for audit archiving](#).

August 30, 2023

This release of the Management Console service introduces the following changes:

Data Lake database upgrade and default major version change

Newly-deployed Data Lake clusters on AWS or GCP with Cloudera Runtime 7.2.7 or above are now configured to use a PostgreSQL version 14 database by default.

Newly-deployed Data Lake clusters on Azure with Cloudera Runtime 7.2.7 or above will continue to use a PostgreSQL version 11 database by default.

The database for Data Lake clusters on AWS and GCP can now be upgraded to PostgreSQL version 14. If your AWS or GCP cluster requires an upgrade to PostgreSQL 14, you will receive a notification in the Management Console UI.

Cloudera strongly recommends that the database upgrade to PostgreSQL 14 for AWS and GCP clusters is performed on all clusters running PostgreSQL version 11 by November 9, 2023.

A database upgrade to PostgreSQL 14 for Azure Data Lakes will be available in the future. Any Data Lake clusters on Azure that require a database upgrade will be upgraded from PostgreSQL 10 to PostgreSQL 11.

For more information, see [Upgrading Data Lake/Data Hub database](#)

August 23, 2023

This release of the Management Console service introduces the following changes:

Configure proxy for TLS interception and Deep Packet Inspection

After setting up a web proxy server in CDP, you can further configure it to perform TLS interception and Deep Packet Inspection (DPI). For instructions, see [Setting up a web proxy for TLS inspection](#).

Setting a default identity provider in CDP

You can optionally set a default identity provider (IdP) in CDP. If you do so, CDP will use the default IdP instead of the oldest IdP. For instructions, see [Setting a default identity provider in CDP](#).

July 20, 2023

This release of the Management Console service introduces the following changes:

Enterprise Data Lakes

A new Data Lake shape called the enterprise Data Lake is now available for new deployments of Cloudera Runtime version 7.2.17. The enterprise Data Lake is a redefined version of medium duty Data Lakes that still offer failure resilience, but utilize resources and allocate memory more efficiently than a medium duty Data Lake at the same cost. For more information see [Data Lake scale](#).

Medium duty Data Lakes deprecated

Medium duty Data Lakes are deprecated as of Runtime 7.2.17. You can upgrade a medium duty Data Lake from 7.2.16 to 7.2.17, but will not be able to upgrade it further. You can create a new 7.2.17 medium duty Data Lake through the CDP CLI, but Cloudera recommends using the enterprise Data Lake for new deployments. The ability to create medium duty Data Lakes will be removed from both the UI and CLI from 7.2.18.

July 19, 2023

This release of the Management Console service introduces the following changes:

CCM enabled by default in the UI

When you register a new AWS, Azure, or GCP environment in CDP via web interface, Cluster Connectivity Manager (CCM) is enabled by default and the option to disable it has been removed from the web interface. Also, CCM is now enabled by default in CDP CLI.

CCMv2 upgrade

CCMv2 upgrade is available for all customers. If your existing CDP environment was created with CCMv1, you will see a notification in your environment details to upgrade to CCMv2. See [Upgrading from CCMv1 to CCMv2](#).

June 28, 2023

This release of the Management Console service introduces the following changes:

Cloudera Runtime 7.2.17

Cloudera Runtime 7.2.17 is now available and can be used for registering an environment with a 7.2.17 Data Lake and creating Data Hub clusters. For more information about the new Runtime version, see [Cloudera Runtime](#). If you need to upgrade your existing CDP environment, refer to [Data Lake upgrade](#) and [Data Hub upgrade](#) documentation.

Rolling Data Lake upgrades

With the release of Cloudera Runtime 7.2.17, you can now upgrade your Data Lake in a rolling manner. The rolling upgrade allows you to upgrade the Data Lake Runtime and OS without stopping attached Data Hubs or Data Services. This allows workloads to continue running during the Data Lake upgrade operation.



Note: You need to contact Cloudera to have this feature enabled.

GCS fine-grained access control (Preview)

You can now register a CDP environment on GCP with RAZ enabled to use fine-grained access policies and audit capabilities available in Apache Ranger. See [GCS Fine-Grained Access Control](#).



Note: You need to contact Cloudera to have this feature enabled.

Cross-version compatibility for Data Lake backup and restore

You can take a backup of a Data Lake that runs one version of Cloudera Runtime and restore the backup to a Data Lake that runs a different version of Runtime. The backup version must be an earlier/lower version Runtime than the Data Lake that you are restoring to. Version limitations apply and a Ranger/HMS schema upgrade may be required. See [Cross-version support for Data Lake backup and restore](#) for more details.

Cross-version compatibility for Data Lakes and Data Hubs

Backward compatibility between Data Lakes and Data Hubs has been introduced with Cloudera Runtime 7.2.17. It is no longer required that you perform Data Hub upgrades in lock-step with the Data Lake upgrade. If your Data Hub is on Runtime version 7.2.16 or later, it is compatible with a Data Lake on a newer Runtime version (7.2.17+). You can independently upgrade your Data Hubs at a later time if you choose to, though it is not required.

AWS Middle East UAE region

You can now register a CDP environment and create Data Hubs in the AWS Middle East UAE region (me-central-1). See updated [Supported AWS regions](#).

Data Lake upgrade validations for Python dependency

If you are planning to upgrade the Runtime version in your existing Data Lake clusters to 7.2.17 or higher versions, you may be required to perform an additional step before upgrading.

You can verify whether your cluster is impacted by navigating to the Upgrade tab of your Data Lake cluster. If there is a warning message about missing prerequisites on the Upgrade page, you need to perform an additional upgrade step before moving to the 7.2.17 Cloudera Runtime version. The required steps may be different depending on your current Runtime version. See [Data Lake upgrade](#) for more details.

June 13, 2023

This release of the Management Console service introduces the following changes:

Zone selection within a GCP region

During GCP environment registration you can now select a zone within the selected GCP region. For example, if you selected the us-west1 region, you can select one of its three zones: us-west1-a, us-west1-b, or us-west1-c.

CDP credit consumption and usage insights

CDP includes a user interface that allows you to monitor your credit consumption and download your consumption records. See [CDP credit consumption and usage insights](#).

April 19, 2023

This release of the Management Console service introduces the following changes:

FreeIPA and Data Lake instance type selection

You can now select the FreeIPA and Data Lake instance types when you register an environment. For more information see [Register an AWS environment from CDP UI](#), [Register an Azure environment from CDP UI](#), and [Register a GCP environment from CDP UI](#).

AWS GP3 support for attached storage disks

AWS Data Lake and FreeIPA instances now support GP3 (SSD) volume types for attached storage. GP3 volumes allow you to increase performance (independently provisioning IOPS and throughput) without increasing storage size. GP3 volumes will deliver similar performance as GP2 volumes at a lower cost. GP3 is now the default attached storage type for instances that previously used GP2 storage.

March 29, 2023

This release of the Management Console service introduces the following changes:

FreeIPA instance type on GCP

Default FreeIPA instance type on GCP changed from n1-standard-2 to e2-standard-2.

January 11, 2023

This release of the Management Console service introduces the following changes:

Cloudera Runtime 7.2.16

Cloudera Runtime 7.2.16 is now available and can be used for registering an environment with a 7.2.16 Data Lake and creating Data Hub clusters. For more information about the new Runtime version, see [Cloudera Runtime](#). If you need to upgrade your existing CDP environment, refer to [Data Lake upgrade](#) and [Data Hub upgrade](#) documentation.

Instance type selection for Data Lake and FreeIPA on AWS

If necessary, you can now select a larger or smaller instance type for a Data Lake or FreeIPA after the environment has been deployed. This is only supported on AWS. For more information see [Vertically scaling instances types - Data Lake](#) and [Vertically scale FreeIPA instances](#).

2022

November 23, 2022

This release of the Management Console service introduces the following changes:

Data Lake resizing

You can now scale up a light duty Data Lake to the medium duty form factor, which has greater resiliency than light duty and can service a larger number of clients. You can trigger the scale-up in the CDP UI or through the CDP CLI. For more information, see [Data Lake resizing](#).

November 3, 2022

This release of the Management Console service introduces the following changes:

Data Lake backup and restore for GCP

Backing up and restoring a GCP Data Lake is now supported. For more information, see [Backup and restore for the Data Lake](#).

October 20, 2022

This release of the Management Console service introduces the following changes:

Azure Reference Network Architecture

New conceptual overview of the CDP Public Cloud network architecture for Azure, its use cases, and personas who should be using it.

For more information, see [Azure Reference Network Architecture](#).

September 27, 2022

This release of the Management Console service introduces the following changes:

Database Upgrade and default major version change

Newly deployed Data Lake and Data Hub clusters with Cloudera Runtime 7.2.7 or above are now configured to use a PostgreSQL version 11 database by default.

A new Database Upgrade capability is now available for existing Data Lake and Data Hub clusters. If you are running clusters on Cloudera Runtime version 7.2.6 or below, upgrade to a more recent version before performing the database upgrade.

The major version of the database used by Data Lake or Data Hub clusters is now also displayed on the Database page of the respective service.

Cloudera strongly recommends that the Database Upgrade is performed on all clusters running PostgreSQL version 10 before November 10, 2022.

For more information, see [Upgrading database to Postgres 11](#)

FreeIPA recipes and recipe type changes

You can register and attach recipes to run on a specific FreeIPA host group. For more information, see [Recipes](#).

The following recipe types have been renamed for Data Hub, Data Lake, and FreeIPA recipes:

- pre-service-deployment (formerly pre-cluster-manager-start)
- post-service-deployment (formerly post-cluster-install)

These changes will not affect existing recipe automation.

September 1, 2022

This release of the Management Console service introduces the following changes:

Validate and prepare for upgrade

Before you perform a Data Hub upgrade, you can run the new Validate and Prepare option to check for any configuration issues and begin the Cloudera Runtime parcel download and distribution. Using the validate and prepare option does not require downtime and makes the maintenance window for an upgrade shorter. For more information see [Preparing for an upgrade](#).

AWS Hong Kong region

You can now register a CDP environment and create Data Hubs in the AWS Hong Kong Region (ap-east-1). See updated [Supported AWS regions](#).

July 28, 2022

This release of the Management Console service introduces the following changes:

Changed permissions for managing proxies in CDP

You no longer need to be a PowerUser to register and manage a proxy in CDP. The new minimal roles are as follows:

- EnvironmentCreator can register a proxy in CDP.
- Owner or SharedResourceUser can view details of a proxy.
- Owner can delete a proxy registration from CDP.

This change has been introduced for new proxy registrations only; That is, proxies registered prior to this change continue to be managed by a PowerUser.

See updated [Setting up a non-transparent proxy in CDP](#).

July 12, 2022

This release of the Management Console service introduces the following changes:

New documentation for CDP Public Cloud upgrade

The CDP Public Cloud upgrade advisor, which gives an overview and FAQ of the upgrade process, is now available. See [CDP Public Cloud upgrade advisor](#).

FreeIPA scaling

You can resize your existing FreeIPA cluster via CDP CLI. Upscaling FreeIPA is recommended after performing Data Lake scaling. For more information, see [Resize FreeIPA](#).

July 1, 2022

This release of the Management Console service introduces the following changes:

Support for Cloudera AI in ap-1 and eu-1 regional Control Planes

Cloudera AI is now supported in the ap-1 (Australia) and eu-1 (Germany) regional Control Planes.

June 29, 2022

This release of the Management Console service introduces the following changes:

New "Advanced Options" section in environment registration wizard

The environment registration UI now features a new "Advanced Options" section on some of the pages, which includes some options that were previously featured in the main UI sections. The options that have been moved to the "Advanced Options" sections include:

- On the Data Access and Data Lake Scaling page:
 - Multi-AZ configuration for Data Lake and FreeIPA (available for AWS only)
 - Recipe selection for Data Lake

More options will be added to the "Advanced Options" in the future.

New option to delete attached volumes during Data Lake repair

When you initiate a repair from the Data Lake Hardware tab, you have the option to delete any volumes attached to the instance. For more information see [Performing manual Data Lake repair](#).

June 27, 2022

This release of the Management Console service introduces the following changes:

Public Endpoint Access Gateway for Azure

During Azure environment registration, you can optionally enable Public Endpoint Access Gateway, which provides secure connectivity to UIs and APIs in Data Lake and Data Hub clusters deployed using private networking, allowing users to access these resources without complex changes to their networking or creating direct connections to cloud provider networks. With this release, Public Endpoint Access Gateway is general availability for AWS and Azure, and it remains preview for GCP. See [Public Endpoint Access Gateway](#).

Generate workload username based on email

By default, workload usernames are generated using the identity provider user ID. Alternatively, you can now generate workload usernames based on users' email addresses. This is useful in cases when the identity provider user ID is an opaque ID, like a uuid or employee ID, which gives equally opaque workload usernames. For more information, see [Generating workload usernames based on email](#).

AWS Jakarta region

You can now register a CDP environment and create Data Hubs in the AWS Jakarta Region (ap-southeast-3). See updated [Supported AWS regions](#).

Support for Operational Database in ap-1 and eu-1 regional Control Planes

Cloudera Operational Database is now supported in the ap-1 (Australia) and eu-1 (Germany) regional Control Planes.

Restricting all Cloudera SSO access

For added security, you can now restrict all Cloudera SSO access (including account administrator access) by contacting Cloudera Support and they can disable or enable the "Cloudera SSO All Login Enabled" setting for the account. Previously, account administrator access could not be restricted. For more information, see [Disabling the Cloudera SSO login](#).

June 16, 2022

This release of the Management Console service introduces the following changes:

Customer managed encryption keys on GCP

By default, a Google-managed encryption key is used to encrypt disks and Cloud SQL instances in Data Lake, FreeIPA, and Data Hub clusters, but you can optionally configure CDP to use a customer-managed encryption key (CMEK) instead. This can only be configured using CDP CLI. There is no UI option available for specifying a GCP CMEK in CDP. For more information, refer to [Adding a customer managed encryption key for GCP](#).

June 7, 2022

This release of the Management Console service introduces the following changes:

Customer managed encryption keys on AWS

By default, Data Lake and FreeIPA's Amazon Elastic Block Store (EBS) volumes and Relational Database Service (RDS) are encrypted using a default key from Amazon's KMS, but you can optionally configure encryption using Customer Managed Keys (CMK). Data Hubs inherit environment's encryption key by default but you have an option to specify a different CMK during Data Hub creation. For more information, refer to [Adding a customer managed encryption key to a CDP environment running on AWS](#).

Deploying CDP in multiple AWS availability zones

By default, CDP provisions Data Lake, FreeIPA, and Data Hubs in a single AWS availability zone (AZ), but you can optionally choose to deploy them across multiple availability zones (multi-AZ). It is possible to enable it either for all or some of these components. For more information, refer to [Deploying CDP in multiple AWS availability zones](#).

June 3, 2022

This release of the Management Console service introduces the following changes:

SCIM for Azure AD

CDP supports SCIM with Microsoft Azure Active Directory (Azure AD). For more information, see [Configure SCIM with Azure AD](#).

May 26, 2022

This release of the Management Console service introduces the following changes:

New permissions were added to the default cross-account AWS policy

The [cross-account access IAM role](#) that is used for the CDP credential has been changed to include a set of new permissions required for Cloudera Data Engineering (CDE), Cloudera DataFlow (CDF), and Cloudera AI. The new AWS permissions are required to simplify the creation of the Kubernetes cluster on AWS. As a result of this change, all customers using or planning to use CDE, CDF, or Cloudera AI in CDP Public Cloud on AWS must update their existing cross-account permissions to ensure that these three data services can be created, enabled, or updated.

If you are using or planning to use CDE, CDF, or Cloudera AI, add the following permissions to the cross-account role:

```
{
  "Effect": "Allow",
  "Action": [
    "ssm:DescribeParameters",
    "ssm:GetParameter",
    "ssm:GetParameters",
    "ssm:GetParameterHistory",
    "ssm:GetParametersByPath"
  ],
  "Resource": [
    "arn:aws:ssm:*:*:parameter/aws/service/eks/optimized-ami/*"
  ]
}
```

May 17, 2022

This release of the Management Console service introduces the following changes:

Data Lake scaling (Preview)

Data Lake scaling is the process of scaling up a light duty Data Lake to the medium duty form factor, which has greater resiliency than light duty and can service a larger number of clients. You can trigger the scale-up in the CDP UI or through the CDP CLI. See [Data Lake scaling \(Preview\)](#).

**Note:**

You need to contact Cloudera to have this feature enabled.

May 12, 2022

This release of the Management Console service introduces the following changes:

Cloudera Runtime 7.2.15

Cloudera Runtime 7.2.15 is now available and can be used for registering an environment with a 7.2.15 Data Lake and creating Data Hub clusters. For more information about the new Runtime version, see [Cloudera Runtime](#). If you need to upgrade your existing CDP environment, refer to [Data Lake upgrade](#) and [Data Hub upgrade](#) documentation.

Support for Replication Manager in ap-1 and eu-1 regional Control Planes

Cloudera Replication Manager is now supported in the ap-1 (Australia) and eu-1 (Germany) regional Control Planes.

May 10, 2022

This release of the Management Console service introduces the following changes:

Bring your own Azure private DNS zone

CDP supports using a private endpoint for Azure Postgres with an existing Azure private DNS zone. The private DNS zone can now be pre-created and provided by you, or created by CDP. Previously, CDP always created the private DNS zone when a private endpoint was created.

See updated [Azure requirements for using a private endpoint for Azure Postgres](#) and [Enabling a private endpoint for Azure Postgres in CDP](#).

Extended upgrade version support for RAZ-enabled environments

Data Lake major/minor version upgrades for RAZ-enabled environments are now available for Runtime versions 7.2.10-7.2.12 to 7.2.14+.

April 19, 2022

This release of the Management Console service introduces the following changes:

Changed FreeIPA Azure VM type

The Azure VM type used for the FreeIPA server was changed from Standard_D3_v2 to Standard_DS3_v2 so that FreeIPA nodes can be encrypted at host. Standard_D3_v2 doesn't support encryption at host while Standard_DS3_v2 does. See updated [Overview of Azure resources used by CDP](#).

Setting IDBroker mappings in a RAZ environment is disabled

If a CDP environment has RAZ enabled, setting IDBroker mappings is disabled during environment creation and when the environment is already running. If your environment has RAZ enabled, you should be using Ranger for authorizing user and group access to the S3 or ADLS Gen 2 cloud storage used by the Data Lake.

Azure Load Balancer

The Standard SKU Azure Load Balancer is used in multiple places in CDP Data Lakes and Data Hubs. It is used as a frontend for Knox in both Data Lakes and Data Hubs, and for Oozie HA in HA Data Hubs. See [Azure Load Balancers in Data Lakes and Data Hubs](#).

Upgrading classic clusters from CCMv1 to CCMv2

You can now upgrade your CDH, HDP, or CDP Private Cloud Base clusters that were previously registered in CDP from CCMv1 to CCMv2. See [Upgrading a classic cluster from CCMv1 to CCMv2](#).

March 29, 2022

This release of the Management Console service introduces the following changes:

Upgrading FreeIPA

To ensure that your FreeIPA nodes are running with the latest patches, you should periodically upgrade your FreeIPA cluster. CDP currently allows you to upgrade all FreeIPA clusters, updating OS-level security patches on the cluster nodes. See [Upgrade FreeIPA](#).

Upgrading the Data Lake

Major/minor version upgrades of Cloudera Runtime and Cloudera Manager are generally available. Data Lake maintenance upgrades for RAZ-enabled environments versions 7.2.7+ are generally available. For more information see [Data Lake upgrade](#).

March 21, 2022

This release of the Management Console service introduces the following changes:

New classic cluster roles

As part of the new authorization model released in 2021, CDP introduces a new account role and resource roles related to classic clusters:

	Roles	Description
New account role	ClassicClustersCreator	This role is required to register a new classic cluster. If this role is not present then the “Add Cluster” button is not visible to the user.
New resource roles	ClassicClusterAdmin ClassicClusterUser	These roles can be assigned on the scope of a specific classic cluster.

For more information, see [Enabling admin and user access to classic clusters](#).

Data Lake backup and restore options

New CLI options have been added to the Data Lake backup and restore feature. These options allow for explicitly including or skipping certain data during a backup and restore operation:

- You can skip or include the backup/restore of the HMS and Ranger databases.
- You can skip or include the HBase Atlas tables, and all Solr collections except ranger_audit.
- You can skip or include the Solr ranger_audit collection.

For more information, see [Configure backups for a Data Lake](#).

March 7, 2022

This release of the Management Console service introduces the following changes:

Public certificate auto-renewal

Most public (Let's Encrypt-issued) certificates for Data Lake and Data Hub clusters will now auto-renew without intervention from a user. For more information, refer to [Managing Certificates](#).

Data Lake recipes

Support for attaching/detaching recipes on a Data Lake cluster is now available through both the CDP UI and CDP CLI. For more information see [Recipes](#).

February 25, 2022

This release of the Management Console service introduces the following changes:

Cloudera Runtime 7.2.14

Cloudera Runtime 7.2.14 is now available and can be used for registering an environment with a 7.2.14 Data Lake and creating Data Hub clusters. See [Cloudera Runtime](#).

Resource list filtering

CDP users other than PowerUsers and CDP administrators can only list the resources that they are authorized to access. Prior to this change, all CDP users were able to list all resources, but if they tried to access a resource that they were not authorized to access, CDP would return an error.

February 18, 2022

This release of the Management Console service introduces the following changes:

Customer managed encryption keys on Azure

By default, local Data Lake, FreeIPA, and Data Hub disks attached to Azure VMs and the PostgreSQL server instance used by the Data Lake and Data Hubs are encrypted with server-side encryption (SSE) using Platform Managed Keys (PMK), but you can optionally configure SSE with Customer Managed Keys (CMK). For more information, refer to [Adding a customer managed encryption key for Azure](#).

February 14, 2022

This release of the Management Console service introduces the following changes:

Support for CDW in ap-1 regional Control Plane

Cloudera Data Warehouse (CDW) is now supported in the ap-1 (Australia) regional Control Plane. To use CDW in this regional Control Planes, your CDP administrator must create a new environment.

January 31, 2022

This release of the Management Console service introduces the following changes:

Workload password policies

In order to bring your workload password complexity requirements in line with company policy, you can set your FreeIPA password policies via CDP web interface and CDP CLI. Password policies can be configured for length, complexity, expiration, and scope. For more information, refer to [Configuring workload password policies](#).

Pull-based audit archiving

Pull-based audit archiving allows you to pull audit events for archiving purposes without any extra configuration beyond Control Plane API usage. For more information refer to [Pull-based audit archiving](#).

Custom images and catalogs

If necessary, you can use a custom Runtime or FreeIPA image for compliance or security reasons. You can then use the CDP CLI to register a custom image catalog and set the custom image within the custom image catalog. For more information refer to [Custom images and catalogs](#).

Support for CDW in eu-1 regional Control Plane

Cloudera Data Warehouse (CDW) is now supported in the eu-1 (Germany) regional Control Plane. To use CDW in one this Control Plane, your CDP administrator must create a new environment.

2021

December 17, 2021

This release of the Management Console service introduces the following changes:

Fine-grained Access Control for ADLS Gen2 and S3

The fine-grained access control for ADLS Gen2 and S3 cloud storage via the Ranger Authorization Service (RAZ) enables Amazon S3 and ADLS Gen2 users to control access per user and per directory in cloud storage. By

specifying Apache Ranger policies for cloud storage, admins can provide home directories and audit capabilities similar to those used with HDFS files in an on-premises or IaaS deployment.

For more information and setup steps refer to:

- [Fine grained access control for AWS](#)
- [Fine grained access control for Azure](#)

Related Information

[Cloudera Management Console Overview](#)

November 22, 2021

This release of the Management Console service introduces the following changes:

AWS Milan region is supported for CDW

Cloudera Data Warehouse (CDW) introduces support for the eu-south-1 (Milan) AWS region. See updated [AWS regions supported by CDP](#).

No-proxy option for non-transparent proxies

When you set up a non-transparent proxy server, you now have the option of configuring specific IP addresses, domains, or subdomains to bypass the proxy. For more information, see [Using a non-transparent proxy](#).

CDW diagnostics collection

You can trigger a diagnostics bundle collection for Cloudera Data Warehouse (CDW). See updated [Send a diagnostic bundle to Cloudera Support](#).

Updated GCP provisioning credential's permissions

A new GCP granular permission is required for creating Data Hubs using the Data Engineering HA template:

```
compute.regionHealthChecks.useReadOnly
```

If your GCP provisioning credential uses a custom IAM role with granular permissions, you should update it to include this permission.

See updated [Service account for the provisioning credential](#).

November 17, 2021

This release of the Management Console service introduces the following changes:

Updated Azure provisioning credential's permissions

The following new Azure permissions are required for the CDP provisioning credential:

```
"Microsoft.Network/loadBalancers/backendAddressPools/join/action",  
"Microsoft.Network/loadBalancers/delete",  
"Microsoft.Network/loadBalancers/read",  
"Microsoft.Network/loadBalancers/write",
```

If you have created a custom role for the CDP provisioning credential, you should update your application registration on Azure, assigning these additional permissions. If you have assigned the built-in Contributor role instead of granular permissions, you do not need to take any action.

Documentation has been updated. See [Prerequisites for the provisioning credential](#).

FreeIPA HA for GCP environments

FreeIPA HA is now supported and used by default for all newly created GCP environments.

November 9, 2021

This release of the Management Console service introduces the following changes:

Cluster Connectivity Manager v2 (CCMv2)

CCMv2 replaces CCMv1. While CCMv1 establishes and uses a tunnel based on the SSH protocol, with CCMv2 the connection is via HTTPS. All new environments created with Runtime 7.2.6 or newer after enabling CCMv2 on your tenant use CCMv2. Existing environments and new environments created with Runtime older than 7.2.6 continue to use CCMv1. All newly registered classic clusters use CCMv2, but previously registered classic clusters continue to use CCMv1. If your CDP tenant has not been granted the CDP_CCM_V2 entitlement yet, it continues to use CCMv1.

The steps to register an environment with CCMv2 are similar to CCMv1 configuration steps. The main differences are:

- If you are deploying in an environment with restricted outbound network access, port 443 needs to be open and new destinations need to be added to the allow list.
- If you are registering a classic cluster, the steps have changed.

For more information, see [Cluster Connectivity Manager](#).

October 26, 2021

This release of the Management Console service introduces the following changes:

Medium duty Data Lakes for GCP

Medium duty Data Lakes for GCP have added an additional gateway node to provide failure resilience for UI and API clients. Load-balanced UI and API access are now available without interruption. For more information see [Data Lake scale](#).

Cloudera Runtime 7.2.12

Cloudera Runtime 7.2.12 is now available and can be used for registering an environment with a 7.2.12 Data Lake and creating Data Hub clusters. See [Cloudera Runtime](#).

September 28, 2021

This release of the Management Console service introduces the following changes:

New GCP permissions for provisioning credential

The list of permissions for the provisioning credential's service account has been updated to include new permission required for load balancing between HA components of the Data Lake. If you are running CDP in GCP, you should update the provisioning credential's service account to include either the Compute Load Balancer Admin ([roles/compute.loadBalancerAdmin](#)) IAM role or the following granular permissions:

- compute.addresses.create
- compute.addresses.delete
- compute.addresses.get
- compute.addresses.use
- compute.instanceGroups.create
- compute.instanceGroups.delete
- compute.instanceGroups.get
- compute.instanceGroups.list
- compute.instanceGroups.update
- compute.instanceGroups.use
- compute.forwardingRules.create
- compute.forwardingRules.delete
- compute.forwardingRules.get
- compute.forwardingRules.list

- `compute.forwardingRules.setLabels`
- `compute.forwardingRules.update`
- `compute.forwardingRules.use`
- `compute.regionBackendServices.create`
- `compute.regionBackendServices.delete`
- `compute.regionBackendServices.get`
- `compute.regionBackendServices.list`
- `compute.regionBackendServices.update`
- `compute.regionBackendServices.use`
- `compute.regionHealthChecks.create`
- `compute.regionHealthChecks.delete`
- `compute.regionHealthChecks.get`
- `compute.regionHealthChecks.list`
- `compute.regionHealthChecks.update`
- `compute.regionHealthChecks.use`

See updated [Permissions for the provisioning credential's service account](#).

September 21, 2021

This release of the Management Console service introduces the following changes:

New authorization model

CDP introduces a new authorization model. The following table summarizes new, changed, and deprecated roles. The roles that are not mentioned in this table are unchanged.

Account roles

	Roles	Description
New account role	<ul style="list-style-type: none"> • <code>EnvironmentCreator</code> 	This is a new account-level role.
Deprecated account roles	<ul style="list-style-type: none"> • <code>EnvironmentAdmin</code> • <code>EnvironmentUser</code> 	These roles have been deprecated in June 2020 and have been removed from the official documentation.

Resource roles

	Roles	Description
New environment resource roles	<ul style="list-style-type: none"> • <code>DataSteward</code> • <code>DataHubCreator</code> 	These roles can be assigned on the scope of a specific environment.
New Data Hub resource role	<ul style="list-style-type: none"> • <code>DataHubAdmin (Technical Preview)</code> 	This role can be assigned on the scope of a specific Data Hub.
New shared resource role	<ul style="list-style-type: none"> • <code>SharedResourceUser</code> 	This role can be assigned on the scope of a specific shared resource (cluster template, credential, image catalog, proxy, or recipe).
New resource role applicable to environments, Data Hubs, shared resources, and classic clusters	<ul style="list-style-type: none"> • <code>Owner</code> 	Grants all permissions required to manage the resource in CDP including the ability to delete it, but does not grant any cluster-level access. The user creating the resource automatically gets the Owner role on that resource.

Steps for assigning roles

- The steps for assigning account roles and managing access to environments are unchanged.
- The steps for managing access to Data Hubs, shared resources (cluster templates, credentials, image catalogs, and recipes), and classic clusters are similar to the steps for managing access to environments: You can use the Manage access option from the resource details page.

Updated documentation

- For updated information about all built-in roles in CDP, refer to [Understanding account-level roles and resource roles](#).
- For updated instructions for how to manage access to resources, refer to [Assigning a resource role to a user](#) and [Assigning a resource role to a group](#).
- Other new and updated documentation:
 - [Enabling admin and user access to environments](#)
 - [Enabling admin and user access to Data Hubs](#)

Dots are now supported in group names

The list of supported characters in group names was extended to include dots. See updated documentation:

- [Creating a group](#)
- [Synchronizing group membership](#)

Improved AWS cloud storage setup documentation

AWS cloud storage setup documentation has been improved to include exact steps for creating the required S3 bucket, IAM policies, and IAM roles. See [Minimal setup for cloud storage](#) and [Onboarding CDP users and groups for cloud storage](#).

September 9, 2021

This release of the Management Console service introduces the following changes:

Cloudera Runtime 7.2.11

Cloudera Runtime 7.2.11 is now available and can be used for registering an environment with a 7.2.11 Data Lake and creating Data Hub clusters. See [Cloudera Runtime](#).

Specifying multiple existing AWS security groups

When using your existing security groups for registering an AWS environment in CDP via CDP CLI, you can provide a comma-separated list of multiple security groups for both Knox (`securityGroupIdForKnox`) and Default (`defaultSecurityGroupId`). This is a CLI-only feature.

Specifying multiple GCP shared subnets

When using an existing shared VPC for registering a GCP environment in CDP via CDP web interface or CLI, you can specify multiple shared subnets.

Support for Bahrain (me-south-1) AWS region

Registering an environment and provisioning Data Hubs is now supported in the Bahrain (me-south-1) AWS region. See [Supported AWS regions](#).

Updated outbound network access destinations

If you are using Cloudera AI, Data Engineering, or DataFlow data services and have restricted egress access, starting on September 7, 2021, you need to add the following new endpoints to your egress rules:

- *.s3.{eu-west-1, us-west-2, ap-southeast-1}.amazonaws.com
- s3-r-w.{eu-west-1, us-west-2, ap-southeast-1}.amazonaws.com
- *.execute-api.{eu-west-1, us-west-2, ap-southeast-1}.amazonaws.com

The region selected should be the region that is geographically closest to where the environment is deployed.

Customers operating in outbound restricted networks will be unable to download docker images, which will impact creating new clusters. Existing environments deployed in outbound restricted networks may experience operational issues, including limited ability to start, scale and repair the data service clusters.

For more information, see:

- [Outbound network access destinations for AWS](#)
- [Outbound network access destinations for Azure](#)
- [Outbound network access destinations for GCP](#)

August 31, 2021

This release of the Management Console service introduces the following changes:

New Shared Resources navigation menu item

Management options for provisioning credentials, proxies, and Data Hub cluster templates, recipes, and image catalogs can now be easily accessed from the new Shared Resources item in the navigation menu.

August 12, 2021

This release of the Management Console service introduces the following changes:

Send a diagnostic bundle to Cloudera support

CDP introduces a web interface for sending a diagnostic bundle to Cloudera support. Currently diagnostics can be collected for Data Lake, FreeIPA, and Data Hub. See [Send a diagnostic bundle to Cloudera support](#).

August 9, 2021

This release of the Management Console service introduces the following changes:

New permission for GCP Logger service account

In addition to the previously documented permissions, if you would like to use a bucket path (gs://<bucket>/<path>) instead of a bucket (gs://<bucket>) for the Logs or Backups bucket, you should assign the storage.objects.list permission to the custom role. See [Minimum setup for cloud storage](#).

"Create public IPs" is disabled with CCM

The Create public IPs option available on the UI during Azure and GCP environment registration is now disabled by default when CCM is enabled.

July 12, 2021

This release of the Management Console service introduces the following changes:

Updated quick starts

The AWS, Azure, and GCP quick starts have been updated to include the optional FreeIPA Backup Location Base introduced in a recent release.

- [AWS quick start](#)
- [Azure quick start](#)
- [GCP quick start](#)

July 8, 2021

This release of the Management Console service introduces the following changes:

FreeIPA HA repair

FreeIPA HA repair is now available for all newly created AWS and Azure environments in CDP. When running in high-availability mode, the identity management system runs multiple instances of FreeIPA on separate hosts. In case

of failure, you can now repair failed hosts using the CDP command-line within one week of a node failing. For more information, see [Repair a FreeIPA instance](#).

"Don't create public IPs" option was renamed

The Don't create public IPs option available during Azure and GCP environment registration was renamed to Create public IPs and is enabled by default.

June 24, 2021

This release of the Management Console service introduces the following changes:

FreeIPA backup location

During AWS, Azure, or GCP environment registration via CDP web interface or CDP CLI, you can optionally specify a separate cloud storage location (FreeIPA "Backup Location Base") for FreeIPA backups. If no separate location is specified, FreeIPA backups are stored in Logs Location Base. For more information, see:

- [Minimal AWS setup for cloud storage](#)
- [Minimal Azure setup for cloud storage](#)
- [Minimal GCP setup for cloud storage](#)

FreeIPA HA enabled by default for new AWS and Azure environments

FreeIPA HA is now enabled for all newly created AWS and Azure environments in CDP. The number of nodes used for the FreeIPA server depends on the Data Lake scale selected: light duty uses 2 nodes and medium duty uses 3. The FreeIPA HA toggle button has been removed from the environment registration UI, but, if needed, it is possible to customize FreeIPA node count when registering an AWS or Azure environment via CDP CLI. For more information, see [Managing FreeIPA](#).

June 21, 2021

This release of the Management Console service introduces the following changes:

Cloudera Runtime 7.2.10

Cloudera Runtime 7.2.10 is now available and can be used for registering an environment with a 7.2.10 Data Lake and creating Data Hub clusters. See [Cloudera Runtime](#).

S3Guard removal

S3Guard is no longer used with newly registered AWS environments using Runtime version 7.2.2 or newer. Consequently, the "Enable S3Guard" environment registration option has been removed and there is no need to create a DynamoDB table for your environment when planning to use Runtime version 7.2.2 or newer. Environments created prior to this change continue to use S3Guard.

June 4, 2021

This release of the Management Console service introduces the following changes:

Updated GCP Quick Start

GCP Quick Start has been updated to include environment registration steps using CDP web interface. See [GCP Quick Start](#).

Updated CCM troubleshooting documentation

CCM troubleshooting documentation has been updated to include information on common cases when connection via CCM fails and steps for collecting information from clusters. See [Troubleshooting CCM](#).

April 29, 2021

This release of the Management Console service introduces the following changes:

Cloudera Runtime 7.2.9

Cloudera Runtime 7.2.9 is now available and can be used for registering an environment with a 7.2.9 Data Lake and creating Data Hub clusters. See [Cloudera Runtime](#).

Registering CDP Private Cloud Base clusters in CDP Public Cloud

You can now register CDP Private Cloud Base clusters as classic clusters in CDP:

- The CDP Private Cloud Base clusters can be registered via Cloudera Manager for use in Replication Manager.
- Additionally, you can register CDP Private Cloud Base clusters via Cloudera Manager and Knox for use in Data Catalog and Replication Manager. This is a technical preview feature that should not be used in a production environment.

For documentation, see [Adding a CDP Private Cloud Base cluster](#).

April 27, 2021

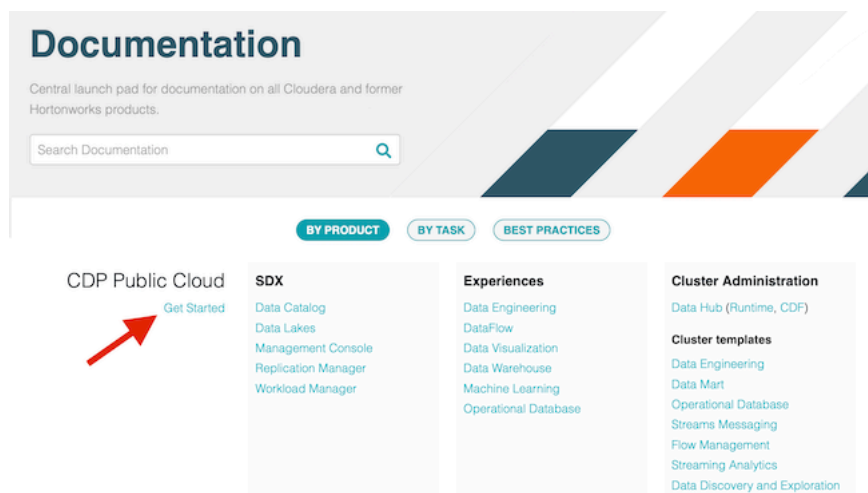
The following updates were made:

CDP Public Cloud onboarding documentation was moved

The following publications were moved to the [CDP Public Cloud](#) library:

- Getting Started in CDP Public Cloud
- AWS/Azure/GCP Requirements
- AWS/Azure/GCP Quick Starts
- CDP Public Cloud Security Overview

This CDP Public Cloud library is accessible via the Get Started link on the <https://docs.cloudera.com> homepage or via the <https://docs.cloudera.com/cdp-public-cloud/cloud/index.html> link:



The documentation that was moved is available from the following links:

- <https://docs.cloudera.com/cdp/latest/getting-started/index.html>
- <https://docs.cloudera.com/cdp/latest/aws-quickstart/index.html>
- <https://docs.cloudera.com/cdp/latest/azure-quickstart/index.html>
- <https://docs.cloudera.com/cdp/latest/gcp-quickstart/index.html>
- <https://docs.cloudera.com/cdp/latest/requirements-aws/index.html>
- <https://docs.cloudera.com/cdp/latest/requirements-azure/index.html>
- <https://docs.cloudera.com/cdp/latest/requirements-gcp/index.html>
- <https://docs.cloudera.com/cdp/latest/security-overview/index.html>

URL redirects were added temporarily; They will eventually be removed, so make sure to update your bookmarks. This change was made in effort to make CDP Public Cloud onboarding documentation easier to find. The previous location of this content (in the Management Console library) was unintuitive to many users.

AWS/Azure/GCP planning documentation was consolidated

The AWS/Azure/GCP requirements content was consolidated in one place in the [CDP Public Cloud](#) library mentioned above.

If the content that you have bookmarked throws a 404 error, it is most likely in one of the following publications:

- <https://docs.cloudera.com/cdp/latest/requirements-aws/index.html>
- <https://docs.cloudera.com/cdp/latest/requirements-azure/index.html>
- <https://docs.cloudera.com/cdp/latest/requirements-gcp/index.html>

To fix the error, you have three options:

- Update the URL by replacing `/management-console/cloud/environments-<cloud-provider>/` with `/cdp-public-cloud/cloud/requirements-<cloud-provider>/`, replacing `<cloud-provider>` with "aws", "azure", or "gcp". This works for the content that was moved, but not for topics that were consolidated into other documentation and removed.
- On the <https://docs.cloudera.com> homepage, search the website for the content that moved. Search results will direct you to the correct location.
- Navigate to one of the libraries linked above and find the content that you are looking for.

This change was made in effort to consolidate all documentation related to cloud provider requirements in one place. Previously, the documentation was scattered and users had to click on many links in order to find content.

April 20, 2021

This release of the Management Console service introduces the following changes:

Cluster Definitions page was moved to environment details

The Cluster Definitions page that used to be available in the Shared Resources section was removed. Instead, you can access all cluster definitions related to a specific environment from the Cluster Definitions tab available in the environment's details. You can save new cluster definitions using the Save As New Definition option available from the Create Data Hub wizard or from CDP CLI using the `cdp datahub create-cluster-definition` command.

Ranger Audit environment parameter was moved to Data Access section

The option to specify the Ranger Audit role (AWS) managed identity (Azure) or service account (GCP) during environment registration was moved from the Logs - Storage and Audit section to the Data Access section. Consequently, these sections were renamed to Logs and Data Access and Audit.

You can select specific nodes to repair within a Data Lake host group

From the Hardware tab of the Data Lake details, you can click the Repair icon to select specific nodes within a host group to repair.

Updated IAM policy for the provisioning credential for AWS

The IAM policy for the provisioning credential has been updated to include new permissions related to load balancers. The following permissions are now required:

```
cloudformation:UpdateStack
cloudformation:ListStackResources
elasticloadbalancing:DescribeLoadBalancers
elasticloadbalancing:DescribeTargetHealth
elasticloadbalancing:RegisterTargets
```



```
elasticloadbalancing:DeregisterTargets
```

If you are using a restricted IAM policy for your provisioning credential, you must add these additional permissions.

April 9, 2020

This release of the Management Console service introduces the following changes:

GCP quick start

GCP quick start is now available, allowing you to quickly set up a CDP environment. See [GCP quick start](#).

March 30, 2021

This release of the Management Console service introduces support for Google Cloud Platform and support for medium duty Data Lakes on Microsoft Azure.

Google Cloud Platform (GCP) support

CDP supports registering GCP environments and creating Data Hub clusters. See the following new and updated documentation:

- [Getting Started as an admin in CDP](#)
- [Google Cloud requirements](#)
- [Introduction to Google Cloud environments](#)
- [Register a Google Cloud environment](#)
- [Medium duty Data Lakes on Google Cloud](#)
- [Default Data Hub configurations available for Google Cloud](#)
- [Create a Data Hub cluster on Google Cloud](#)

Medium duty Data Lakes on Microsoft Azure

The medium duty Data Lake configuration is now available for Microsoft Azure. Light duty is still used by default, but you can change this when registering an environment from CDP user interface or when creating a Data Lake from CDP CLI using the `--scale MEDIUM_DUTY_HA` option. For information about available configurations, see [Data Lake scale](#).

March 25, 2021

This release of the Management Console service introduces the following changes:

Cloudera Runtime 7.2.8

Cloudera Runtime 7.2.8 is now available and can be used for registering an environment with a 7.2.8 Data Lake and creating Data Hub clusters. See [Cloudera Runtime](#).

Medium duty Data Lake for AWS

The medium duty Data Lake configuration is now available for AWS. Light duty is still used by default, but you can change this when registering an environment from CDP user interface or when creating a Data Lake from CDP CLI using the `--scale MEDIUM_DUTY_HA` option. For information about available configurations, see [Data Lake scale](#).

New supported AWS regions

The eu-south-1 (Europe - Milan) and af-south-1 (Africa - Cape Town) regions are now available as technical preview for creating Data Hub clusters. See updated [Supported AWS regions](#).

Cluster definitions location

Cluster definitions that can be used with a given environment are now listed in the details of that environment. To view them navigate your environment and access the Cluster Definitions tab.

March 19, 2021

This release of the Management Console service introduces the following changes:

Single existing Azure resource group

The option to use your existing Azure resource group is now available in CDP. This allows you to have your credential's role definition scoped to that particular resource group instead of the whole subscription. The option to "Select Resource Group" is available on the UI on the Region, Networking and Security page of the register environment wizard. The corresponding CLI JSON parameter is `resourceGroupName` and the `cdp` environment's `create-azure-environment` CLI option to enable is `--resource-group-name <value>`. If these parameters are not present, CDP defaults to creating new resource groups. See updated [Azure Permissions](#) and [Resource group](#) in Azure Requirements documentation, and updated [Register an Azure environment](#).

Private endpoints for Azure database for PostgreSQL

The option to create private endpoints instead of service endpoints for Azure Database for PostgreSQL is now available when registering an Azure environment in CDP. The option to "Create Private Endpoints" is available on the UI in the "Network" section of the register environment wizard. The corresponding CLI JSON parameter is `createPrivateEndpoints` and the `cdp` environments `create-azure-environment` CLI option to enable private endpoints is `--create-private-endpoints`. The option is disabled by default. It can only be used with the single resource group feature and can only be enabled on subnets that have Azure network policies turned off. See [Private endpoint for PostgreSQL](#).

Public Endpoint Access Gateway for AWS

During AWS environment registration, you can optionally enable Public Endpoint Access Gateway, which provides secure connectivity to UIs and APIs in Data Lake and Data Hub clusters deployed using private networking, allowing users to access these resources without complex changes to their networking or creating direct connections to cloud provider networks. See [Public Endpoint Access Gateway](#).

Generate CLI command from the UI

You can quickly obtain CDP CLI commands from the CDP web interface for creating the following:

- Environment: From details of an existing environment or from the last page of the register environment wizard. See [Obtain CLI commands for registering an environment](#).
- Data Lake: From details of an existing Data Lake. See [Understanding Data Lake details: Show CLI Command](#).
- Credential: From details of an existing credential or from the create credential wizard. See [Obtain CLI commands for creating a credential](#).
- Data Hub: From details of an existing Data Hub or from the create Data Hub wizard. See [Create a Data Hub cluster on AWS from an existing cluster](#) and [Create a Data Hub cluster on Azure from an existing cluster](#).

Disable cloud storage logging for an existing environment

By default, CDP sends collected service logs from VM nodes to the cloud storage that you provided for logs during environment registration. You can disable cloud storage logging for a specific environment, by navigating to environment details > Summary > Telemetry and turning off "Enable Cloud Storage Logging". Disabling this option will affect only newly created Data Hub clusters in that environment. See [Enabling environment telemetry](#)>Disable cloud storage logging for an existing environment.

Obtain your CDP tenant ID

You can now obtain your CDP tenant ID from CDP web interface,. See [Obtain your CDP tenant ID](#).

February 25, 2021

This release of the Management Console service introduces the following new features and behavioral changes:

Cloudera Runtime 7.2.7

Cloudera Runtime 7.2.7 is now available and can be used for registering an environment with a 7.2.7 Data Lake and creating Data Hub clusters. See [Cloudera Runtime](#).

User delete

CDP administrators now have the ability to delete users in CDP through both the user interface and the CLI. Deleting a user removes all access keys and SSH keys associated with the user, and unassigns all roles and resource roles assigned to the user. The user is also removed from all groups that they belong to. For more information, refer to [Deleting users and machine users](#).

FreeIPA HA

CDP administrators can configure your CDP environment to run FreeIPA in high-availability mode. By default, creating an environment configures a single instance of FreeIPA on its own host, but you can explicitly enable FreeIPA HA during environment registration via CPD web UI or CLI. For more information, refer to [Managing FreeIPA](#).

Interactive login for CDP CLI and CDP SDK

If you would prefer that user access to the CLI/SDK is shorter-lived, you can use the "interactive" method of logging into the CDP CLI/SDK. By default, this login method grants a 12-hour access key to the CLI/SDK. The access key will time out after one hour of inactivity. The interactive method integrates with any SAML-compliant external identity provider. For more information, refer to [Logging into the CDP CLI/SDK](#).

Anonymization rules

CDP includes a set of default anonymization rules and allows CDP administrators to define custom anonymization rules in order to remove sensitive information from CDP logs. For more information, refer to [Defining anonymization rules for CDP logs](#).

Changes to delete machine user behavior

Deleting a machine user removes all access keys and SSH keys associated with the machine user, and unassigns all roles and resource roles assigned to the machine user. The machine user is also removed from all groups that they belong to. Previously, these steps had to be performed manually prior to machine user deletion. It takes around 2 minutes to fully delete a machine user in CDP. During that time you will not be able to recreate the machine user (that is, for 2 minutes you will not be able to create a machine user with the same machine user name).

Group name length limit

CDP user management framework supports group names of up to 64 characters. Previously up to 32 characters were supported.

Identity provider configuration improvements

The user interface and the overall flow of the identity provider configuration in CDP was improved for better usability.

New CDP SAML Service Provider certificate

The current CDP SAML Service Provider certificate is expiring on March 8, 2021 at 18:05:49 GMT. A replacement certificate is available for any customer whose identity provider will verify the CDP SAML service provider certificate. You can obtain the certificate from this document or by logging it to CDP web interface, navigating to > User Management > Identity Providers, clicking on your identity provider, and the last field "CDP SAML Service Provider Metadata" now contains 2 certificates: the one that expires on March 8, 2021 and the new one. Please consult your identity provider's documentation for how to update service provider certificates. CDP will start using the new certificate for SAML starting March 8, 2021.

Here is the new CDP SAML Service Provider certificate:

```
-----BEGIN CERTIFICATE-----
MIIEKTCCAxEgAwIBAgIU7LjObY+L8dcCVzWN4ChnTtybiowDQYJKoZIhvcNAQEL
BQAwgAMxCzAJBgNVBAYTAlVTMQswCQYDVQQIDAJDQTEWMBQGA1UEBwwNU2FuIEZy
YW5jaXNjbzEVMBMGA1UECgwMQ2xvdWRlcmEgSW5jMRAwDgYDVQQLDAdDRFAGSUFN
MSEwHwYDVQQDDbhjb25zb2x1LmNkcC5jbG91ZGVyYS5jb20xIzAhBgkqhkiG9w0B
CQEFWFHN1CHBvcnRAY2xvdWRlcmEgSW5jMRAwDgYDVQQLDAdDRFAGSUFNMSEwHw
YDVQQDDbhjb25zb2x1LmNkcC5jbG91ZGVyYS5jb20xIzAhBgkqhkiG9w0BCQEF
WFHN1CHBvcnRAY2xvdWRlcmEgSW5jMRAwDgYDVQQLDAdDRFAGSUFNMSEwHwYDV
QQDDbhjb25zb2x1LmNkcC5jbG91ZGVyYS5jb20xIzAhBgkqhkiG9w0BAQEF
AAOCAQ8AMIIBCgKCAQEAzIxvwxsAE4PqNlfZ2+4zfYI9UpiePEOKJuLlQ8Mbh
ArA53EmZradpYNIQ54a3vGQNeEoi782gcp/JbzLTY0AESnKXzpPXOhX8qMWytrcL
QKmSW/eVbZsVEYnyf1wFxtPocLbHfYB12W1ScD3FKr5BUms6bRCc1fiFW1Ei5XLQ
yzgSGdKXSvB/8izRr4yyyDT2IX8PelHbONiIKb6OTuuHPwo259RMjZZd2pwMurif
JUGBckwYPh7Dkmiw9mTXVSD5fdSP1HvP2RTuPqmkTSgJRwdJD4G6wF1NFOQWitIr
7vf6OzPZJM6A2JCN8RQApMnYyNgT75wWtCNOf8F2cQIDAQABo1MwUTAdBgNVHQ4E
FgQUGFVSdXrVb3JsJy5nf4OYp2sJn8IwHwYDVR0jBBgwFoAUGFVSdXrVb3JsJy5n
f4OYp2sJn8IwDwYDVR0TAQH/BAUwAwEB/zANBgkqhkiG9w0BAQsFAAOCQAQEAkNxx
+X2sCbXAIhSUNKUYQEM++ZDSnWzMGdavnVUzWgTfGdwvDolFzvqU66wiQ8kedK0
qLW6gRZkG+GUJq5vY93pfNSQ5C4P9hhFqpd6tfHme7uHlZCtZh/wjOeYoOpgr0eI
qtXxg6U6+6qLqzBi/9Zdc0sLZFNbjQLEfkNHOU7rFODcnLNHemngw+ui2rofsBhK
F9ZcqiY9lmmCto6OrQMAkXyfrU40S8+Yr9s+wnJEmNikVN9mfH0TfRJNEvHcuvZ+
WHc4HD/Vu0sL/APPADfLh158MYb9gUNXtE12PxjGYCj4RsFt0/Fbju9mGl+W/n69
qaRFxZmubutaQ1WCzw==
-----END CERTIFICATE-----
```

CDP CLI reference

CDP CLI reference documentation is now available at <https://cloudera.github.io/cdp-dev-docs/cli-docs/>.

Documentation for configure lifecycle management for logs on AWS and Azur

To avoid unnecessary costs related to Amazon S3 or ADLS Gen2 cloud storage, you should create lifecycle management rules for your cloud storage location used by CDP for storing logs so that these logs get deleted once they are no longer useful. See [Configure lifecycle management for logs on AWS](#) and [Configure lifecycle management for logs on Azure](#).

Consolidated documentation for restricting admin and end user access for CDP services

Consolidated documentation for restricting admin and end user access for CDP services is now available. Previously these options were only covered in documentation related to specific CDP workload services. See [Restricting access for CDP services that create their own security groups on AWS](#) and [Restricting access for CDP services that create their own security groups on Azure](#).

Updated AWS and Azure requirements documentation

AWS and Azure requirements documentation was updated to include more requirements related to Data Engineering, Data Warehouse, and Cloudera AI. These requirements were previously only documented in service-specific docs. See updated [AWS Requirements](#) and [Azure Requirements](#).

January 19, 2021

This release of the Management Console service introduces the following new features and behavioral changes:

Control Plane audit archiving

You can configure Control Plane audit archiving from CDP web interface. Previously, this feature was only available via CDP CLI. For updated documentation, refer to [AWS setup for audit archiving](#) and [Azure setup for audit archiving](#).

New documentation for pre-creating ADLS Gen2 account for storing OS images

CDP uses an ADLS Gen2 storage account for storing images used for VMs. By default, CDP creates this account during environment registration, but you can optionally pre-create it. If needed, you can also copy the VHD files and create image resources manually. For instructions on how to do this, refer to [ADLS Gen2 account for storing images](#).

January 6, 2021

This release of the Management Console service introduces the following new features and behavioral changes:

Python 3.6 is required for CDP CLI

Starting in January 2021, the CDP CLI requires Python 3.6 for execution. [Python 2 reached end of life on January 1, 2020](#) and is no longer receiving updates. While your existing CDP CLI installation running on Python 2 will keep working, new features and bug fixes are available only for CDP CLI installations running on Python 3.6. To check your Python version, use:

```
python -V
```

If you need to install Python 3, you can:

- Visit <https://www.python.org/downloads/>
- Use a package manager for your OS
- Use a Python management tool

Once you have Python 3, install CDP client by following the usual [CDP client installation instructions](#).

2020

December 21, 2020

This release of the Management Console service introduces the following new features and behavioral changes:

New functionality

- Data Lake metadata backup and restore: You can backup and restore the metadata maintained in the Data Lake services. Data Lake backup and restore is supported from Cloudera Runtime 7.2.1+ on AWS and Cloudera Runtime 7.2.2+ on Azure. See [Backup and restore for the Data Lake](#).

Changed functionality

- The option to set tenant-level tags was moved from Environments > Shared Resources > Tags to a new location under Global Settings > Tags.
- The option to enable workload analytics and deployments cluster logs collection on a tenant-level was moved from Environments > Shared Resources > Telemetry to a new location under Global Settings > Telemetry.

New documentation

- [Azure resources used by CDP](#)

Updated documentation

- [Azure permissions](#)
- [Renew host certificates on Data Lake and Data Hub clusters](#)

Deprecated functionality

The following account roles are deprecated:

- EnvironmentAdmin
- EnvironmentUser

November 20, 2020

This release of the Management Console service introduces the following new features and behavioral changes:

New documentation

- [AWS outbound network access destinations](#)
- [Azure outbound network access destinations](#)
- [Reserved group names](#)
- [VM-based on demand diagnostics](#) (technical preview)

Updated documentation

- Updated requirements for [Azure Database for PostgreSQL](#)
- IAM role and managed identity setup for access to cloud storage was updated to grant the IDBROKER_ROLE or Assumer identity the same permissions as the LOG_ROLE or Logger identity, allowing IDBroker to write to logs storage location. See updated [Minimal setup for cloud storage](#) (AWS) and [Minimal setup for cloud storage](#) (Azure).
- CDP CLI, API, and SDK documentation was moved out of the Management Console library and is now available from the general CDP Public Cloud library. See [CDP APIs](#), [CDP CLI](#) and [CDP SDK](#).

October 27, 2020

This release of the Management Console service introduces the following new features and behavioral changes:

New features

- [Renew host certificates on Data Lake and Data Hub clusters](#)

New documentation

- [Prerequisites for adding classic clusters with a non-transparent proxy](#)
- [Creating a machine user in CDP](#)

Updated documentation

- [Register an AWS environment](#) - Reflects updated environment wizard
- [Register an Azure environment](#) - Reflects updated environment wizard

September 24, 2020

This release of the Management Console service introduces the following new features and behavioral changes:

New features

- [cdpcurl utility](#): The cdpcurl utility has been released to [open source](#). This provides curl-like access to the [Management Console API](#) as another alternative to the CDP CLI.

Behavioral changes

- The option to stop and restart a Data Lake was removed. You should stop and restart your whole environment instead, as described in [Stop and restart an environment](#).
- The option to delete a Data Lake was removed from the Actions menu on the environment details page. When deleting an environment from the CDP UI, you no longer need to delete the environment's Data Lake separately. See updated [Delete an environment](#) documentation.
- The option to Check for Data Lake Upgrade was removed.
- The options to Enable Workload Analytics and Enable Cluster Logs Collection were removed from the Actions menu on the environments details page. You can now control these options from the Management Console > Environments > Shared Resources > Telemetry (tenant-level setting), or you can set them during environment

registration (environment-level setting). For updated instructions, see [Enable workload analytics and cluster logs collection](#).

- The option to Get FreeIPA certificate were removed from the environment's Actions menu and is now available from FreeIPA's Actions menu. To find the option, navigate to the Management Console > Environments > navigate to a specific environment > Summary > FreeIPA > Actions > Get FreeIPA certificate.

New documentation

- [Performing user sync](#)
- [Configuring Azure Active Directory identity federation in CDP](#)
- [VNet and subnet planning](#) for Azure
- [Supported AWS regions](#) includes an updated list of AWS regions where Cloudera AI can be deployed
- [Supported Azure regions](#) includes an updated list of Azure regions where Cloudera AI can be deployed
- [Default security group settings](#) for Azure

August 24, 2020

This release of the Management Console service introduces the following new features:

SSH key management

A Power User can add and delete SSH keys for all users and users can add and delete their own SSH keys. For more information, see [Managing SSH keys](#).

Choose Data Lake version at deployment

When deploying an environment, you can now select the Data Lake version. See updated instructions for [Register an AWS environment](#) and [Register an Azure environment](#).

Edit parameters of a previously deployed environment

You can edit the following environment components:

- Add new subnets: See [Add subnets to an environment](#).
- Add new security groups: See [Add security groups to an environment](#).
- Add a new root SSH key: See [Add root SSH key to an environment](#).

July 31, 2020

This release of the Management Console service introduces the following new features:

CDP Control Plane Public API Reference

CDP Control Plane now features reference information for APIs located at [CDP Control Plane Public API Documentation](#). You can use this reference documentation to look up descriptions and syntax for SDK functions and CLI commands.

June 30, 2020

This release of the Management Console service introduces the following new features:

Groups sync disabled by default

When you create a Group in CDP, the Sync Membership flag is now off by default. This will eliminate the situation when users would be removed from a CDP group if they weren't part of the same group in the ID Provider.

Support for HDP 2.6.x Classic Clusters using CCM

Support for registering Classic clusters using the reverse SSH mechanism is extended to legacy HDP 2.6.x clusters. This will be behind the CLASSIC_CLUSTERS entitlement.

Auditing Control Plane activity

Auditing is used to collect or log evidence of activity in a system that auditors can use to both track and analyze to answer questions such as: Who made a change to the system? When did a change happen? What exactly changed? Why was a change authorized?

May 30, 2020

This release of the Management Console service introduces the following new features:

Edit subnets

You can now add new subnets to the configuration of a running environment. This might be useful when you add a new service such as Cloudera AI and Cloudera Data Warehouse (CDW) which require their own subnet.

Improvements in Data Hub cluster scaling

A number of significant improvements were made to properly clean up cloud resources for scaling up and down Data Hub clusters. This will reduce cloud costs that would have been incurred by any orphaned instances.

April 27, 2020

This release of the Management Console service introduces the following new features:

Support for connecting to private subnets

CDP supports communicating with Data Lake and Data Hub workload clusters that are on private subnets. environments in Azure.

You can use Cluster Connectivity Manager (CCM) to communicate with Data Lake and Data Hub workload clusters that are on private subnets. Communication takes place over private IPs without any inbound network access rules required. CDP requires that these clusters have outbound connections to AWS NLBs hosted in Cloudera's AWS account. Workload clusters initiate an SSH tunnel to the CDP control plane, which is then used for all communication thereafter.

Proxy server (technical preview)

Proxy server works with DataLake and Datahub in both transparent and non-transparent mode. The proxy server works with Cloudera AI and Cloudera Data Warehouse (CDW) in transparent mode only. Proxy server does not currently work with Cloudera AI and CDW in non-transparent proxy mode.

You can use a proxy server to control the connections that are allowed from your VPC or VNet and prevent unattended connections initiated from your environment.

April 6, 2020

This release of the Management Console service introduces the following new features:

Tagging CDP resources

When you create an environment or other resources shared across your cloud provider account, CDP automatically adds default tags to the Cloudera-created resources in your cloud provider account. You can also define additional custom tags at either the environment level or tenant level.

February 28, 2020

This release of the Management Console service introduces the following new features:

Creating environments in Azure

CDP supports creating environments in Azure.

Register your Azure account and then launch Data Hub clusters, Data Warehouse clusters, and Cloudera AI workbenches within the Azure environment and determine which users have access to which resources. These

clusters are attached to a Data Lake that runs within the environment and provides security and governance for all attached clusters.

2019

December 19, 2019

This release of the Management Console service introduces the following new features:

Specifying multiple CIDRs on security groups

CDP supports specifying multiple comma-separated CIDRs during environment registration under Security Access Settings > Access CIDR.

Enabling workload analytics for Data Hub clusters

For each environment, you can manually enable and disable workload analytics so that diagnostic information about job and query execution is sent to Workload Manager for Data Hub clusters created within this environment. The option is available during environment creation under Logs Storage and Audits > Enable Workload Analytics. You can also update it once the environment is running by navigating to environment details > Actions > Enable/Disable Workload Analytics:

IAM role selection

When providing IAM instance profiles or IAM roles required for environment's Logs Storage and Audits and Data Access configuration, you can now select from available instance profiles and roles instead of manually providing IAM role ARNs.

November 14, 2019

This release of the Management Console service introduces the following new features:

Wire encryption

CDP now offers TLS encryption across all endpoints for both data and control traffic, providing an internal Certificate Authority (CA) with commonly trusted certificates and automatic certificate expiry warnings. Adding new cluster hosts or services to a cluster with auto-TLS enabled automatically (which is CDP's default setting) creates and deploys the required certificates. CDP supports TLS 1.2.

Existing DynamoDB table

When configuring S3Guard as part of environment creation, you can now specify an existing DynamoDB table. Note that:

- The table must be in the same region as the environment
- One table can be used with one environment only
- CDP will not delete this table when deleting the environment

Removal of environment groups

There is no longer need to create the `cdp_$env` group for each newly created environment and assign it to the environment.

Consistent FreeIPA password across all environments

There is no longer need to reset the FreeIPA password after creating a new environment. The password is automatically propagated to each newly created environment. You can still manage your FreeIPA password from the User Management page and reset it if needed.

Knox uses port 443 instead of 8443

Knox gateway now uses port 443. If you select for CDP to create security groups automatically, this port is automatically open to your organization's CIDR as needed. If you are creating your environment's security groups manually, you should open port 443 instead of 8443 to your organization's CIDR.

September 23, 2019

This release of the Management Console service introduces the following new features:

Adding or updating IDBroker mappings on a running environment

To add or update the IDBroker mappings on a running environment, navigate to Environments > select an environment > Actions > Manage Access > IDBroker Mappings. For more information, refer to [Editing IAM role to S3 mappings](#).

Retry option for Data Lake

When stack provisioning or Data Lake cluster creation fails, the retry option to resume the process from the last failed step. For more information, refer to [Retry a Data Lake](#).

August 22, 2019

This is the first release of the Management Console service. For an overview of Management Console functionality, refer to:

Kubernetes Ingress NGINX Controller vulnerabilities

Five vulnerabilities affecting the Ingress Nginx Controller for Kubernetes were publicly disclosed on March 24, 2025, and were given the nickname 'IngressNightmare'.

The 'IngressNightmare' vulnerabilities may allow Remote Code Execution (RCE) and potentially expose Kubernetes clusters to malicious configuration modifications. Exploitation requires specially crafted HTTP requests that bypass security measures, such as a Web Application Firewall (WAF). Successful exploitation may lead to complete cluster compromise, data exfiltration, and denial of service.

Details of the CVEs:

- [CVE-2025-1974](#) (CVSS score: 9.8) – An unauthenticated attacker with access to the pod network can achieve arbitrary code execution in the context of the ingress-nginx controller under certain conditions
- [CVE-2025-24514](#) (CVSS score: 8.8) – The auth-url Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- [CVE-2025-1097](#) (CVSS score: 8.8) – The auth-tls-match-cn Ingress annotation can be used to inject configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- [CVE-2025-1098](#) (CVSS score: 8.8) – The mirror-target and mirror-host Ingress annotations can be used to inject arbitrary configuration into NGINX, resulting in arbitrary code execution in the context of the ingress-nginx controller and disclosure of secrets accessible to the controller
- [CVE-2025-24513](#) (CVSS score: 4.8) – An improper input validation vulnerability that could result in directory traversal within the container, leading to denial-of-service (DoS) or limited disclosure of secret objects from the cluster when combined with other vulnerabilities

How does these vulnerabilities affect Cloudera on cloud?

Cloudera Control Plane is a Software as a Service platform and is not vulnerable as the admission controller is disabled on all control planes. The Cloudera Control Plane will be patched to remediate this issue by March 31, 2025.

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2025-839: Critical Kubernetes Ingress NGINX Controller Vulnerability Allows RCE Without Authentication](#)

Log4j vulnerabilities

Cloudera has released hotfixes for Public Cloud Runtime versions 7.2.7 and newer that address Log4j2 vulnerabilities.

The following vulnerabilities have been addressed for Public Cloud Runtime versions 7.2.7 through 7.2.12:

[CVE-2021-44228](#)

[CVE-2021-45046](#)

[LOGBACK-1591](#)

[CVE-2021-45105](#)

[CVE-2021-44832](#)

You should upgrade your CDP services running Runtime versions 7.2.7+ so that they include the latest hotfixes. You can update your existing Data Lake and Data Hubs by doing a maintenance upgrade. You should first upgrade the Data Lake and then upgrade all the Data Hubs that are using the Data Lake. Refer to [Data Lake upgrade](#) and [Data Hub upgrade](#) documentation. Data Lake and Data Hub maintenance upgrade is supported only in technical preview for maintenance upgrades from Runtime versions 7.2.7 and higher.

If you are running a version of Runtime prior to 7.2.7, contact Cloudera Support for details on how to upgrade Runtime.

For more information about these hotfixes, refer to the [CDP Public Cloud Runtime Release Notes](#) for the version of Runtime that you use.

Image catalog updates

This section lists prewarmed image updates used to Data Lakes, Cloudera Data Hub clusters and FreeIPA.

The following prewarmed images were added on the following dates:

First released	Platform	Cloudera Runtime component versions	Cloudera Manager version	Available OSs
April 3, 2024	Cloudera Runtime 7.2.18	7.2.18 Component Versions	Cloudera Manager 7.12.0	RHEL 8.8 EUS
June 27, 2023	Cloudera Runtime 7.2.17	7.2.17 Component Versions	Cloudera Manager 7.11.0	RHEL 8.8 EUS, CentOS 7
January 11, 2023	Cloudera Runtime 7.2.16	7.2.16 Component Versions	Cloudera Manager 7.9.0	CentOS 7
May 12, 2022	Cloudera Runtime 7.2.15	7.2.15 Component Versions	Cloudera Manager 7.6.2	CentOS 7
February 25, 2022	Cloudera Runtime 7.2.14	7.2.14 Component Versions	Cloudera Manager 7.6.0	CentOS 7
October 25, 2021	Cloudera Runtime 7.2.12	7.2.12 Component Versions	Cloudera Manager 7.5.2	CentOS 7
September 9, 2021	Cloudera Runtime 7.2.11	7.2.11 Component Versions	Cloudera Manager 7.4.3	CentOS 7
June 14, 2021	Cloudera Runtime 7.2.10	7.2.10 Component Versions	Cloudera Manager 7.4.2	CentOS 7

First released	Platform	Cloudera Runtime component versions	Cloudera Manager version	Available OSs
April 30, 2021	Cloudera Runtime 7.2.9	7.2.9 Component Versions	Cloudera Manager 7.4.1	CentOS 7
March 25, 2021	Cloudera Runtime 7.2.8	7.2.8 Component Versions	Cloudera Manager 7.4.0	CentOS 7
February 9, 2021	Cloudera Runtime 7.2.7	7.2.7 Component Versions	Cloudera Manager 7.3.0	CentOS 7
December 11, 2020	Cloudera Runtime 7.2.6	7.2.6 Component Versions	Cloudera Manager 7.2.6	CentOS 7
October 8, 2020	Cloudera Runtime 7.2.2	7.2.2 Component Versions	Cloudera Manager 7.2.2	CentOS 7
July 31, 2020	Cloudera Runtime 7.2.1	7.2.1 Component Versions	Cloudera Manager 7.2.1	CentOS 7
June 16, 2020	Cloudera Runtime 7.2.0	7.2.0 Component Versions	Cloudera Manager 7.2.0	CentOS 7
February 20, 2020	Cloudera Runtime 7.1.0	7.1.0 Component Versions	Cloudera Manager 7.1.0	CentOS 7

For the latest image catalog updates, refer to Cloudera web interface > Cloudera Management Console > Shared Resources > Image Catalogs.

Known issues for Cloudera Management Console

This section lists known issues that you might run into while using the Cloudera Management Console service.

Data Lake upgrade fails during backup with out of memory error

Problem: When upgrading the Data Lake from Cloudera Runtime 7.2.18 to one of its service pack versions or to 7.3.1, the upgrade fails with the following out of memory error during backup:

```
Upgrade not started, datalake backup failed. Failure message: Database: kinit: Client 'hdfs/xxxxxxxxxx.CLOUDERA.SITE' not found in Kerberos database while getting initial credentials; kinit: Client 'hbase/xxxxxxxxxx.CLOUDERA.SITE' not found in Kerberos database while getting initial credentials; kinit: Client 'solr/uxxxxxxxxxxx.CLOUDERA.SITE' not found in Kerberos database while getting initial credentials; moveFromLocal: Failed with java.io.IOException while processing file/directory :[/xxxxxxxx-xxx-xxxx-xxxx-xxxxxxxf_database_backup/ranger_backup._COPYING_] in method:[java.util.concurrent.ExecutionException: java.lang.OutOfMemoryError: Direct buffer memory];
```

Workaround: You need to manually increase the value of Client Java Heap Size in Cloudera Manager using the following steps:

1. Navigate to your environment in Cloudera Management Console, and click on the Data Lake tab on the environment details page.
2. Open Cloudera Manager by clicking on the CM URL.
3. Select Clusters core_settings in Cloudera Manager.
4. Click on the Configuration tab.
5. Search for Client Java Heap Size in Bytes in the search bar.
6. Increase the value from the default 256 MB to 1 GB (for example, 1073741824 bytes).
7. Click Save changes.
8. Restart the affected services as prompted by Cloudera Manager to apply the new configuration.

After the affected services are started again, retry the upgrade operation. Monitor the upgrade process to confirm the issue is resolved and the Data Lake upgrade finishes without errors.

CDPD-50436 Kudu service user not authorized to access Hive warehouse in RAZ enabled AWS cluster

Problem: The Kudu service user is not authorized to access Hive warehouse locations in a RAZ enabled AWS cluster on cloud object stores, which under certain conditions can prevent Kudu tables from being created. This results in the following error:

```
ImpalaRuntimeException: Error creating Kudu table 'default.truckspeedevents'
CAUSED BY: NonRecoverableException: failed to create HMS catalog entry for
table [id=b764bceb167746b7bb3dc1e8722e66e6]: failed to create Hive MetaStor
e table: TException - service has thrown: MetaException(message=Got exceptio
n: java.nio.file.AccessDeniedException
```

Workaround: Add "kudu" service user to the allow list for "Default: Hive warehouse locations" in the cm_s3 Ranger repository in your S3 cloud storage.

ENGESC-19520 Cloudera can't validate permissions

Problem: If in your AWS organization you have Service Control Policies (SCPs) where certain regions are enabled or disabled, Cloudera can't validate the permissions correctly due to an AWS SDK limitation.

Workaround: Disable the permission validations.

OPSAPS-64580 installation is failing because of failure in Install Longhorn step

Problem: When using Cloudera Manager to install , the installation may fail due to a failure in the Longhorn install step. You will see an error similar to the following in the stderr output:

```
++ openssl passwd -stdin -apr1 + echo 'cm-longhorn:$apr1$gp2nrbtq$1KYPGIOQN1
FJ2lo5sV62l0' + kubectl -n longhorn-system create secret generic basic-auth
--from-file=auth + rm -f auth + kubectl -n longhorn-system apply -f /opt/clo
udera/cm-agent/service/ecs/longhorn-ingress.yaml Error from server (Internal
Error): error when creating "/opt/cloudera/cm-agent/service/ecs/longhorn-ing
ress.yaml": Internal error occurred: failed calling webhook "validate.nginx.
ingress.kubernetes.io": Post "https://rke2-ingress-nginx-controller-admissio
n.kube-system.svc:443/networking/v1/ingresses?timeout=10s": x509: certificat
e signed by unknown authority
```

Workaround: Retry the installation by clicking the Resume button.

OPSX-735 Kerberos service may report errors if Cloudera Manager is not running

Problem: The Cloudera Manager Server in the base cluster must be running in order to generate Kerberos principals for Cloudera on premises. If there is downtime, you may observe Kerberos-related errors.

Workaround: Resolve downtime on Cloudera Manager. If you encountered Kerberos errors, you can retry the operation (such as retrying creation of the Virtual Warehouse).

OPSX-1405 Multiple users creating the same environment may result in an unusable environment

Problem: If two users try to create an environment with the same name at the same time, it might result in an unusable environment.

Workaround: Delete the existing environment and try again with only one user creating the environment.

OPSX-1412 Creating an Environment may fail

When multiple users try to create the same environment at the same time or use automation to create an environment with retries, environment creation may fail on collision with a previous request to create an environment.

Workaround: Delete the existing environment, wait 5 minutes, and try again.

OPSX-2062 Platform not shown on the Compute Cluster tab in Environments

Problem: In the Cloudera Management Console Environments page the platform does not display on the Compute Clusters tab.

OPSX-2713 ECS Installation fails to perform First Run of services.

Problem: If an issue is encountered during the Install Cloudera Control Plane step of Containerized Cluster First Run, installation will be re-attempted infinitely rather than the command failing.

Workaround: Since the control plane is installed and uninstalled in a continuous cycle, it is often possible to address the cause of the failure while the command is still running, at which point the next attempted installation should succeed. If this is not successful, abort the First Run command, delete the Containerized Cluster, address the cause of the failure, and retry from the beginning of the Add Cluster wizard. Any nodes that are re-used must be cleaned before re-attempting installation.

OPSX-2772 Unable to modify permissions for Account Administrator

Problem: When a user with administrative privileges accesses the User Management Update Roles page in the Cloudera Management Console, the user is presented with options to select various roles. Selecting or deselecting these roles does not change this user's privileges -- an administrative user, by default, has all privileges, and those privileges cannot be changed.

CB-11925: Knox Load Balancer API requests initiated from Knox Gateway hosts can fail with Connection timeout error

Problem: When logged into a Data Lake or Cloudera Data Hub node that has a Knox Gateway service instance configured on it, making Knox API calls through the Knox load balancer can result in a connection timeout error. This is because for security reasons, the IP address of the request is preserved in the traffic passed through the load balancer. Preserving the IP address means that the load balancer will reject "loopback" traffic, meaning traffic that originates and is directed back to the same node.

Workaround: If Knox API calls need to be made while logged into a Knox gateway node, use the hostname of the node instead of the load balancer hostname in the API call.

The Knox load balancer hostname can be identified by the "-gateway" suffix in the first clause of the hostname with no numeric identifier. For example: <cluster-name>-gateway.<env-shortname>.<hash>.cloudera.site is the load balancer hostname, and <cluster-name>-gateway0.<env-shortname>.<hash>.cloudera.site is a direct node hostname.

OPSAPS-59129: Cloudera Manager reports the HDFS warning "Secure DataNode configuration is valid, but not recommended."

Problem: On Cloudera Runtime versions 7.2.9 and earlier, Data Lake clusters that include the HDFS service show this warning in Cloudera Manager: "Secure DataNode configuration is valid, but not recommended." This warning is benign and can be ignored.

CDPSDX-2879 Ranger import fails when you create a Hive replication policy for a medium duty Data Lake cluster

Problem: When you create a Hive replication policy with the Include Sentry Permissions with Metadata or Skip URI Privileges option for a medium duty Data Lake cluster, Ranger import fails. Before you choose the Include Sentry Permissions with Metadata option for a Hive replication policy for a medium duty Data Lake cluster, contact Cloudera Support.

CB-10535, CB-10372 Generate CLI command for existing environment should show 3 commands instead of 1

Problem: If you try to obtain the CDP CLI commands from an existing environment > Actions > Show CLI commands, only the create environment command is displayed instead of all three commands required for registering an environment from CDP CLI.

Workaround: You can obtain the command for creating a Data Lake from Data Lake details. The command to obtain the set IDbroker mappings can be obtained from an existing environment or from CDP CLI help, but you need to modify it manually to set the same mappings as in the source environment.

CB-10706 SSO is not working for Solr/Namenode UI links

Problem: SSO login to an environment with a medium duty Data Lake breaks access to Solr and Namenode UI links.

Workaround: After you deploy a medium duty Data Lake, login to Gateway0 and run:

```
openssl rand -base64 12
```

Then login as root to both gateway nodes and run:

```
export KNOX_GATEWAY_DATA_DIR=/var/lib/knox/gateway/data
```

```
/opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh create-alias pac4j.password --cluster knoxsso --value "the value from above"
```

Then in Cloudera Manager, restart the Knox service.

DWX-6635 Tags are not being added to S3 buckets

Problem: S3 buckets that are part of your AWS environment registered in Cloudera are not being tagged during environment registration. This is because the PutBucketTagging policy is missing from the cross-account IAM role that Cloudera requires you to create for your environment's credential.

Workaround: You can:

- Manually add tags to your S3 buckets used for existing environments.
- Add the PutBucketTagging policy to the IAM role used for your provisioning credential so that any environments registered in Cloudera in the future can automatically add S3 bucket tags.

CB-6924 Workaround for ZooKeeper external volume bug

Problem: In the current version of Cloudera, ZooKeeper might be configured to write to Cloudera's root disk which is too small to accommodate the ZooKeeper data. To correct this issue, you need to reconfigure ZooKeeper to write to an external volume and move any ZooKeeper data to that volume.

Workaround: To check if ZooKeeper is configured to use an external volume, complete the following:

1. Open ZooKeeper and navigate to: ZooKeeper menu item -> Configuration tab -> Filter to Server.
2. If the dataDir and dataLogDir fields contain /hadoopfs/fs1/zookeeper you do not need to do anything.
3. If the fields contain any other values, you must reconfigure ZooKeeper.

To reconfigure ZooKeeper, complete the following:

1. ssh into the machine where the ZooKeeper server is running .
2. Run the following command to change the user:

```
sudo -su zookeeper
```

3. Run the following command:

```
cp -R /var/lib/zookeeper/ /hadoopfs/fs1/zookeeper
```

4. Open the cluster from the Cloudbreak user interface.
5. Log into the Cloudera Manager user interface.
6. Find ZooKeeper on the Cloudera Manager page and navigate to the configuration with either the Search box or select it from the side menu: ZooKeeper menu item -> Configuration tab -> Filter to Server.

7. Change the following properties:

```
dataDir: /hadoopfs/fs1/zookeeper
dataLogDir: /hadoopfs/fs1/zookeeper
```

8. Save your changes
9. Restart the Stale configuration

You do not need to redeploy ZooKeeper.

CB-3876 Cloudera Data Warehouse and Cloudera AI create security groups

Problem: When during environment registration you choose to use your own security groups, the Cloudera Data Warehouse and Cloudera AI services do not use these security groups but create their own.

Workaround: For instructions on how to restrict access on the security groups created by the Cloudera Data Warehouse service, refer to [Restricting access to endpoints in AWS environments](#).

CRB-971 Cloudera Data Warehouse creates IAM, S3, and DynamoDB resources

Problem: The Cloudera Data Warehouse service creates its own S3 buckets, DynamoDB tables, and IAM roles and policies. It does not use the environment's S3 bucket(s), DynamoDB table, and IAM roles and policies.

Workaround: There is no workaround.

CB-4176 Data Lake cluster repair fails after manual stop

Problem: Data Lake cluster repair fails after an instance has been stopped manually via AWS console or AWS CLI.

Workaround: After stopping a cluster instance manually, restart it manually via the AWS console or AWS CLI, and then use the Sync option in Cloudera to sync instance state.

CB-2813 Environment with Cloudera AI workbenches in it can be deleted

Problem: When deleting an environment that uses a customer-created VPC and subnets, there is no mechanism in place to check for any existing Cloudera AI workbenches running within the environment. As a result, an environment can be deleted when Cloudera AI workbenches are currently running in it.

Workaround: If using an environment created within your existing VPC and subnets, prior to deleting an environment, ensure that there are no Cloudera AI workbenches running within the environment.

CB-3459 Subnet dependency error when deleting an environment

Problem: When deleting an environment that uses a VPC and subnets created by Cloudera, the environment deletion fails with an error similar to:

```
com.sequenceiq.cloudbreak.cloud.exception.CloudConnectorException: AWS Cloud
Formation stack reached an error state: DELETE_FAILED reason: The subnet 'su
bnet-05606fd72fda58c8c' has dependencies and cannot be deleted. (Service: Am
azonEC2; Status Code: 400; Error Code: DependencyViolation; Request ID: da9a
7fe0-ac43-467e-9942-94f10e6bd2b7)
```

This error occurs if there are resources such as instances used for Cloudera Data Warehouse, or Cloudera AI cluster nodes that were not deleted prior to environment termination.

Workaround: Prior to terminating an environment, you must terminate all clusters running within that environment.

CB-4248 Expired certificate causes untrusted connection warning

Problem: Cloudera automatically generates an SSL certificate for every Data Lake and Cloudera Data Hub cluster. There are two possibilities:

- By default, Cloudera generates a trusted certificate valid for 3 months.
- If generating a trusted certificate fails, Cloudera generates a self-signed certificate valid for 2 years.

In the first case, if your cluster stays active for over 3 months, the trusted certificate will expire and you will see an "untrusted connection", warning when trying to access cluster UIs from your browser.

Workaround: To fix this, you should generate a new certificate by using the following steps:

1. Use the Renew certificate UI option:

- For Data Lake - Click the Renew certificate button on the Data Lake details page.
- For Cloudera Data Hub - Click Actions > Renew certificate on Cloudera Data Hub cluster details page.

During certificate renewal, several related messages will be written to Event History. Once the certificate renewal has been completed, the following message appears: "Renewal of the cluster's certificate finished."

2. Additionally, if your cluster was created prior to December 19, you need to perform the following manual steps:

- a. SSH to the Knox gateway host on your cluster.
- b. Run the hostname command to get your domain name.
- c. Run the following commands (just replace the domain name test-master.dev.cldr.work with your correct, fully-qualified domain name):

```
sudo sh -c '/opt/salt_2017.7.5/bin/salt --out=newline_values_only 'test-master.dev.cldr.work' pillar.get gateway:userfacingcert > /etc/certs-user-facing/server.pem'
sudo systemctl reload nginx.service
```

CB-23926: Autoscaling must be stopped before performing FreeIPA upgrade

Problem: During autoscaling, compute nodes are stopped and cannot receive the latest FreeIPA IP address update in a stopped state. This causes the Cloudera Data Hub clusters to become unhealthy after the FreeIPA upgrade when autoscaling is enabled.

Workaround:

1. Disable autoscaling before performing FreeIPA upgrade.
2. Start all compute nodes, which are in a STOPPED state.
3. Perform the FreeIPA upgrade.
4. Enable autoscaling after FreeIPA upgrade is successful.

DOCS-27053: Edit SCP policy when all regions except where network resources are created are blocked

Error: When provisioning a new AWS environment, customers with strict Service Control Policy (SCP) rules that block all regions except where their network resources are created may encounter the following error.

Error message:

```
You are not authorized to perform this operation. User: arn:aws:sts::assumed-role/CDP-Cross-Account_Cert_ROLE/cdp-provisioning is not authorized to perform : ec2:DescribeRouteTables with an explicit deny in a service control policy (Service: Ec2, Status Code: 403, Request ID:) Request ID: URL:
```

This happens because, immediately after credential selection during environment registration, a network resource query is initiated. This query is performed via an API call that defaults to the us-east-1 region. If the customer's network resources are located in another region, for example, eu-south-1, and an SCP policy blocks all regions except eu-south-1, then the error is likely to be encountered.

Solution: If the customer is using an existing Cloudera credential with permission verification disabled and have later implemented an SCP policy that blocks the us-east-1 region, then they need to edit and save the credential (Cloudera Console > Shared Resources > Credentials > edit and save the credential) to handle this error message automatically.

Behavioral changes in Cloudera Management Console

Behavioral changes denote a marked change in behavior from the previously released version to Cloudera Management Console service.

Summary:

Cloudera environments require more secure SSH keys

Previous behavior:

Cloudera Data Hub clusters could be created without any failure using older SSH keys.

New behavior:

Due to security improvements, Cloudera environments require BC FIPS SSH keys. When the Cloudera environment does not have an SSH key with sufficient strength, the Cloudera Data Hub cluster provisioning fails. In this case, you need to generate an SSH key that complies with BC FIPS, and replace the existing SSH key with the newly generated one in the Cloudera environment. After updating the SSH key, Cloudera recommends rotating the Cloudbreak user root SSH public key for FreeIPA, Data Lake and the Cloudera Data Hub clusters.

You can update and rotate the SSH key using the following steps:

For UI steps

1. Navigate to Cloudera Management Console, and your environment.
2. Select **Summary** tab, and scroll down to the **Advanced** section.

The screenshot shows the Cloudera Management Console interface. On the left is a sidebar with navigation links. The main content area displays the 'Summary' tab for an environment. Below the summary, the 'Advanced' section is visible, showing the 'Root SSH:' field. The 'Event History' panel on the right lists several events, including 'Environment creation successfully finished' and 'FreeIPA creation/registration started for environment'.

3. Click Edit, and provide the New SSH public key.

The screenshot shows the 'Advanced' section of the Cloudera Management Console. A warning message states: 'The newly added ssh key will only be used for the Cloudera Data Hub clusters created within the environment after the new ssh key were added. All the existing environment resources such as the Data Lake, FreeIPA, and any existing Cloudera Data Hub clusters will remain within the ssh key originally defined during environment creation.' Below the warning, the 'Root SSH:' section has two radio buttons: 'New SSH public key' (selected) and 'Existing SSH public key'. A text input field contains the placeholder '<new-ssh-key>'. At the bottom are 'Save' and 'Cancel' buttons.

4. Click Save.

After saving the changes, you can rotate the Cloudbreak user root SSH public key for FreeIPA, Data Lake and the Cloudera Data Hub clusters. For more information, see the following documentations:

- [Rotating FreeIPA secrets](#)
- [Rotating Data Lake secrets](#)
- [Rotating Cloudera Data Hub secrets](#)

For CLI steps

Use the following CLI command to update the SSH key for the Cloudera environment.

```
cdp environments update-ssh-key --environment [***ENVIRONMENT  
NAME***] --new-public-key [***NEW SSH KEY***]
```

After updating the SSH key in CDP CLI, you can rotate the Cloudbreak user root SSH public key for FreeIPA, Data Lake and the Cloudera Data Hub clusters. For more information, see the following documentations:

- [Rotating FreeIPA secrets](#)
- [Rotating Data Lake secrets](#)
- [Rotating Cloudera Data Hub secrets](#)